Equity Prices the Opportunity, Debt Prices the Risk: Generative AI and Corporate Credit Spreads.*

PRIYANK GANDHI, JUNTAI LU, JASPER PAN, ALBERTO PLAZZI AND JIA WEI†

This version: May 12, 2026

Abstract

How do corporate bond markets price workforce exposure to generative AI? We find that firms with greater workforce exposure to generative AI experience credit spreads nearly 8 basis points wider per standard deviation of exposure following the release of ChatGPT, well exceeding the 4 – 6 basis point effects for policy uncertainty shocks documented in the literature. This finding appears paradoxical: high-exposure firms simultaneously experience rising equity valuations and wider credit spreads. We argue that ChatGPT's release created a real option to restructure labor whose fundamental parameters — adoption costs, success probabilities, and competitive consequences — were unknown at the outset. This parameter uncertainty, typically absent from prior models of labor-technology substitution, simultaneously raises expected asset value and generates relatively wider credit spreads for high-exposure firms. A model with parameter uncertainty rationalizes this pattern, and cross-sectional tests confirm its predictions: the effect is larger for firms with greater credit fragility, labor intensity, and weaker governance.

JEL Codes: G12, G14, G32, G33, G34, J23, O33.
Keywords: Generative artificial intelligence, credit spreads, real options, labor-technology substitution

*We thank participants at the Eastern Finance Association for comments and suggestions. All errors are our responsibility.

†Gandhi is with Rutgers Business School (corresponding author), E-mail: priyank.gandhi@business.rutgers.edu. Lu is with Auburn University at Montgomery, E-mail: jlu4@aum.edu. Pan is with The College of New Jersey, E-mail: jasper.pan@rutgers.edu. Plazzi is with the Swiss Finance Institute and Università della Svizzera italiana (USI Lugano), E-mail: alberto.plazzi@usi.ch. Wei is with Kennesaw State University, E-mail: jwei4@kennesaw.edu.

Introduction

Generative artificial intelligence (AI) is reshaping the global economy at a pace and scale that few technologies in recent history can match. How financial markets price this disruption and how investors in firms exposed to generative AI fare are critical empirical questions. Using the public release of ChatGPT on November 30, 2022 as an information event, we analyze credit spreads on a large panel of U.S. corporate bonds. We find that firms with greater workforce exposure to generative AI experienced credit spreads that are higher by nearly 8 basis points per standard deviation of exposure relative to otherwise similar low-exposure firms following ChatGPT's release. This magnitude exceeds the 4 to 6 basis point effects documented for policy uncertainty shocks in credit markets (Kaviani, Kryzanowski, Maleki, and Savor, 2020), consistent with generative AI representing a more pervasive source of firm-level uncertainty than episodic policy shocks. Scaled to the nearly \$1.74 trillion bonds outstanding for high-exposure firms, the 8 basis point difference in borrowing costs implies approximately \$1.39 billion in additional annual interest, with implications for corporate investment, capital structure, and the real economy.

It is not ex ante obvious that corporate bond spreads should be higher for high-exposure firms. Existing literature offers two channels with competing predictions for how generative AI affects corporate bond spreads. The first, which we call the *disruption channel*, following Becker and Ivashina (2023), predicts that technological disruptions erode the market share and profitability of incumbent firms, resulting in higher default rates and wider credit spreads. Kogan, Papanikolaou, and Stoffman (2020) formalize this logic in a general equilibrium model, and Knesl (2023) provides direct empirical evidence, documenting that risk premia increase for firms with high shares of displaceable labor when automation technology improves. Applied to our setting, this channel predicts that high-exposure firms, i.e., the incumbents most threatened by automation, should face higher borrowing costs following ChatGPT's release relative to firms with low exposure.

The second, which we call the *substitution channel*, follows Zhang (2019) and predicts that credit spreads will tighten. Zhang (2019) demonstrates that firms with a large share of labor engaged in routine tasks hold a real option to substitute technology for workers, thereby hedging firm value against macroeconomic shocks and reducing systematic risk. If generative AI strengthens this substitution mechanism, high-exposure firms should become safer for bondholders, causing credit spreads to contract. The prediction aligns with the aggregate evidence in Andrews and Farboodi (2025), who document that long-term corporate bond yields fall following major AI model releases, consistent with markets pricing the long-run productivity gains that AI enables. Eisfeldt, Schubert, and Zhang (2026) provide corroborating evidence from the cross-section of equity returns: valuations for high-exposure firms increased following ChatGPT's release. Under standard models of credit risk, this implies tighter, rather than wider, credit spreads.

Our results show that the disruption channel dominates for corporate bonds: relative to otherwise similar low-exposure firms, high-exposure firms face borrowing costs that are significantly higher following ChatGPT's release, an effect also observed for subsequent major AI model launches. This finding, however, presents a puzzle. The firms that face relatively higher borrowing costs are precisely the same firms whose equity valuations rose, as documented by Eisfeldt, Schubert, and Zhang (2026), since we employ identical measures of firm-level generative AI exposure. Consequently, shareholders of high-exposure firms gained even as the same firms faced relatively higher borrowing costs. Shareholders and bondholders, reacting to the same technology shock within the same firms, reached opposite conclusions about how generative AI affects firm risk. This divergence — same firms, same event, opposite reactions — warrants closer scrutiny.

We argue that the resolution of this puzzle lies in the nature of the real option that generative AI creates for high-exposure firms. The release of ChatGPT and subsequent major AI model launches provide high-exposure firms with a credible opportunity to substitute technology for labor and to restructure their workforces at a scale previously unfeasible. Evidence that firms are already exercising this option is clear: technology sector layoffs increased from approximately 95,000 in 2024 to over 170,000 in 2025, with AI-driven restructuring cited as a primary motivation. In addition, postings for entry-level positions most exposed to AI fell 35% between 2023 and 2025.[1] These patterns confirm that generative AI represents a genuine, large-scale restructuring opportunity for

[1]Layoff statistics are from Crunchbase (2025) and Layoffs.fyi (2025). The decline in entry level jobs is from Simon (2025) as reported by CNBC (2025).

high-exposure firms.

The real option created by generative AI exerts two opposing forces on firm value. First, generative AI increases expected firm value. Successful workforce restructuring enables high-exposure firms to perform the same volume of work with fewer employees, thereby reducing labor costs, improving productivity, and generating higher expected cash flows. This is precisely the mechanism equity investors value: Eisfeldt, Schubert, and Zhang (2026) document that market valuations rose sharply for high-exposure firms following ChatGPT's release, a finding consistent with equity markets pricing in the expected productivity gains enabled by successful generative AI-led restructuring.

The same real option, however, also increases uncertainty. Significant ambiguity remains regarding how high-exposure firms will navigate the generative AI transition: the success and timing of workforce transitions, the associated restructuring costs, and the resulting shifts in product-market competition are all uncertain, and a positive outcome is by no means guaranteed. While some firms will successfully exercise the option, becoming leaner and more profitable, others will stumble, incurring significant losses or facing displacement by more innovative competitors. Even for successful firms, the potential increase in free cash flow introduces an agency problem: shareholders may extract these gains through dividends and share buybacks rather than allowing them to accrue to bondholders. Anticipating higher uncertainty and potential agency costs, bondholders rationally price them into credit spreads, causing high-exposure firms to face relatively higher borrowing costs.

In the standard structural credit risk framework of Merton (1974), shocks — such as improvements in generative AI — that widen the distribution of possible firm outcomes increase asset volatility, simultaneously raising equity value and increasing credit spreads relative to unaffected firms. This is because equity (as a call option) is a convex claim on firm assets that benefits from increased variance, while debt is a concave claim (a short put option) that loses value when downside risk rises. Shareholders price the option value that generative AI creates; bondholders price the uncertainty surrounding its exercise. The same technology therefore generates opposing signals across the capital structure, not because equity and debt investors disagree about AI's potential, but because their claims have fundamentally different payoff structures.

To formalize this intuition, we develop a simple theoretical framework building on Merton (1974) and Zhang (2019). In the model, improvements in generative AI create a real option for high-exposure firms to substitute technology for labor. Unlike the substitution option in Zhang (2019), whose parameters are known in equilibrium and whose optimal exercise truncates the left tail of the asset value distribution, benefiting both shareholders and bondholders, the generative AI option arrives with unknown fundamental parameters (such as its adoption cost, success probability, or competitive consequences). This parameter uncertainty widens the distribution of possible outcomes in both directions, raising both expected asset value and asset volatility. Expected asset value rises because successful restructuring reduces operating costs and improves expected cash flows. Asset volatility rises because unknown parameters prevent optimal exercise (Miao and Wang, 2011), fattening the left tail of the asset value distribution rather than truncating it. In the Merton (1974) framework, these two forces operate asymmetrically across the capital structure: the simultaneous increase in expected asset value and volatility raises equity value and generates relatively wider credit spreads for high-exposure firms.

Our framework clarifies why the disruption channel dominates the substitution channel in our setting. In Zhang (2019), the labor-technology substitution option pre-exists, its parameters are known in equilibrium, and it is exercised countercyclically as a hedge against adverse aggregate shocks, reducing firm systematic risk, raising equity value, and tightening credit spreads. In our setting, the option is newly created, its parameters are unknown, and it is triggered by technological innovation rather than by adverse aggregate conditions. This parameter uncertainty, rather than mere timing uncertainty about a known option, is what generates asset volatility, determines which channel dominates, and produces the equity-debt divergence we document.

We take the model to the data and confirm its testable predictions. First, equity values rise for high-exposure firms while their credit spreads are significantly wider relative to low-exposure firms — shareholders gain even as firms face relatively higher borrowing costs, an effect that scales with the degree of workforce exposure to generative AI. Second, implied equity volatility and analyst forecast dispersion both rise significantly for high-exposure firms, confirming that option markets and financial analysts immediately incorporate the parameter uncertainty surrounding the AI restructuring option into their pricing and forecasts. Third, the increase in implied volatility is asymmetric and concentrated in the left tail of the firm value distribution, as reflected in higher deep out-of-the-money put implied volatility for high-exposure firms. Fourth, since an increase in asset volatility has a larger impact on credit spreads as firms approach default, the spread widening is stronger for firms with greater credit fragility — proxied by credit ratings, Altman Z-scores, equity volatility, and leverage. Finally, all results are moderated for firms with strong corporate governance — specifically, those with higher institutional ownership and greater CEO pay sensitivity to equity performance — consistent with bond investors viewing well-governed firms as less likely to exercise the restructuring option opportunistically.

Overall, our results are robust to several changes in empirical design. These include alternative fixed-effects structures that can absorb unobserved firm heterogeneity, and alternative methodologies to measure credit spreads. The latter includes using data from credit default swap (CDS) markets, volume-weighted bond yields, within-bond standardized spreads, and value-weighted yields of portfolios sorted on generative AI exposure. We confirm that pre-period coefficients are uniformly small and statistically insignificant, thus supporting the parallel trends assumption. In addition, Fisher randomization inference, which randomly reassigns generative AI exposure across firms 1,000 times, places our estimated coefficient at the 100^{th} percentile of the permutation distribution, rejecting the null of no effect at the 0.1% level. A pre-period placebo test, which assigns a placebo date one year prior to ChatGPT's release, produces a near-zero and statistically insignificant coefficient. Our results are also robust to controlling for time-varying idiosyncratic equity volatility. Following Campbell and Taksler (2003), controlling for idiosyncratic equity volatility or augmenting the baseline specification with Post $\times$ Volatility interaction leaves the AI exposure coefficient stable, confirming that the effect is not mechanically driven by any general increase in equity volatility following generative AI improvements.

Our work makes three key contributions to the literature. First, we extend and bridge the frameworks of Zhang (2019) and Becker and Ivashina (2023), demonstrating that the economic channels in their respective work are not independent but reflect two sides of the same real option, one pricing the potential upside and the other pricing the inherent uncertainty. We derive conditions under which each channel dominates. In our setting, where the option is newly created and carries substantial parameter uncertainty, the disruption channel prevails.

Second, we provide the first cross-sectional analysis of how the risks and opportunities associated with generative AI are priced in corporate bond markets. While seminal studies focus on equity markets (Eisfeldt, Schubert, and Zhang, 2026; Babina, Fedyk, He, and Hodson, 2024) or aggregate bond yields (Andrews and Farboodi, 2025), the cross-section of corporate debt has remained largely unexplored. We document that high-exposure firms simultaneously experience rising equity valuations and relatively higher credit spreads than low-exposure firms — a divergence that persists across successive AI model launches. We show this follows directly from the concave structure of debt claims and the parameter uncertainty surrounding the exercise of the real option that generative AI creates.

Finally, we connect the labor economics literature on AI-driven occupational displacement (Eloundou, Manning, Mishkin, and Rock, 2024; Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025) to the pricing of corporate debt. We establish that the workforce disruption documented in those studies has measurable and economically significant consequences for corporate borrowing costs, transforming labor market shifts into tangible financial market risk premia.

The remainder of this paper proceeds as follows. Section 1 reviews the related literature. Section 2 describes the data and presents summary statistics. Section 3 presents the baseline results and robustness tests. Section 4 develops the theoretical framework and derives testable predictions. Section 5 takes the model to the data. Finally, Section 6 concludes.

1 Related literature

This paper sits at the intersection of five literatures: innovative disruption and financial markets, labor-technology substitution and asset pricing, the determinants of corporate credit spreads, parameter uncertainty in asset pricing, and real options and corporate finance.

Prior work establishes that innovative disruption impacts firms asymmetrically, creating winners and losers (Papanikolaou, 2011; Kogan and Papanikolaou, 2014; Kogan, Papanikolaou, Seru, and Stoffman, 2017; Kogan, Papanikolaou, and Stoffman, 2020). Becker and Ivashina (2023) show that in debt markets, disruption widens credit spreads for incumbent firms. Recent work analyzes the impact of improvements in generative AI. Eisfeldt, Schubert, and Zhang (2026) find that firms with high-AI-exposure earn higher equity returns, Andrews and Farboodi (2025) document that aggregate bond yields fall following major AI releases, and Babina, Fedyk, He, and Hodson (2024, 2025) show that AI adoption reshapes firm productivity and industry concentration.[2] We are the first to examine how generative AI impacts the cross-section of bond spreads. Our results reveal that while aggregate yields decline — likely reflecting expected economy-wide productivity gains — firm-specific disruption risk generates relatively wider spreads for high-exposure incumbents — a divergence invisible in aggregate data alone.

A second strand examines how innovation impacts labor markets and asset prices. Kogan, Papanikolaou, Schmidt, and Seegmiller (2023) classify innovations as labor-saving or labor-augmenting, with the former more likely to generate disruption risk. Eloundou, Manning, Mishkin, and Rock (2024) classify generative AI as labor-saving and Hampole, Papanikolaou, Schmidt, and Seegmiller (2025) confirm that labor demand falls as firm AI exposure rises. Zhang (2019) — the paper most closely related to ours — shows that firms with high routine-task labor shares hold a real option to substitute technology for workers, which hedges firm value and lowers expected returns. Knesl (2023) challenges this finding — firms with high displaceable labor earn 4% higher returns. Our paper reconciles these studies and identifies the conditions under which the substitution channel breaks down and disruption risk prevails, generating relatively wider credit spreads

[2] Relatedly, Andreadis, Chatzikonstantinou, Kalotychou, Louca, and Makridis (2025) find municipal bond yields decline for counties with greater AI labor investment post-ChatGPT.

for high-exposure firms.

Our work also builds on the structural credit risk model of Merton (1974). We extend this framework by showing that generative AI introduces a new source of asset volatility — workforce restructuring uncertainty — which is priced in credit markets. In addition, following the literature, throughout we control for aggregate credit conditions and market liquidity (Collin-Dufresne, Goldstein, and Martin, 2001; Chen, Lesmond, and Wei, 2007; Gilchrist and Zakrajšek, 2012). Coles, Daniel, and Naveen (2006) and Brockman, Martin, and Unlu (2010) show that governance dampens credit spreads, and we test for this using several governance proxies. Finally, Campbell and Taksler (2003) establish that idiosyncratic equity volatility is a significant driver of bond yields, motivating our test of whether the AI shock raised implied volatility and whether this increase in volatility accounts for wider spreads for high-exposure firms.

A fourth strand examines how parameter uncertainty affects asset prices (Epstein and Schneider, 2010). Miao and Wang (2011) demonstrate that parameter uncertainty widens the distribution of outcomes and increases asset volatility. Drechsler (2013) models uncertainty as a stochastic process that steepens the volatility smirk in options markets. Chen and Han (2023) confirm empirically that options market uncertainty widens credit spreads, particularly for low-rated firms. We apply this framework to ChatGPT's release, showing that the generative AI shock produces higher asset volatility that bondholders price into credit spreads, with the left tail of asset value distribution reflected in deep OTM puts, consistent with Drechsler (2013).

Finally, our paper relates to the real options literature. McDonald and Siegel (1986) and Dixit and Pindyck (1994) establish that investment irreversibility creates an option to wait that fundamentally alters firm behavior. In our framework, ChatGPT's release creates a real option to restructure the workforce whose parameters are unknown, generating uncertainty on top of standard investment timing risk. The results in Novy-Marx (2011), Donangelo (2014), and Serfling (2016) show that labor restructuring costs amplify firm risk, directly motivating our finding that the spread differential is largest for labor-intensive firms.

2 Data and summary statistics

Our sample integrates corporate bond transaction data from TRACE and FISD with firm-level generative AI workforce exposure scores from Eisfeldt, Schubert, and Zhang (2026), augmented with firm-level data from Compustat and CRSP. Detailed variable definitions appear in Table A1 in Appendix A1.

2.1 Bond market data

We construct a panel of U.S. corporate bonds over 2021–2024 — a two-year window on either side of ChatGPT's public release on November 30, 2022. Bond transaction data come from FINRA's TRACE, combining the Enhanced, Standard, and Rule 144A datasets to maximize coverage. We follow Dick-Nielsen (2009) and Dick-Nielsen (2014) to clean the raw data, retaining non-convertible, non-asset-backed, fixed-rate coupon bonds with a maturity of at least 1 year denominated in U.S. dollars, and applying the decimal-shift and bounce-back filters of Dickerson, Robotti, and Rossetti (2023). We aggregate the cleaned data to daily bond-level observations, computing volume-weighted average clean price across all transactions for each bond on each day.

For each bond-day observation, we compute the yield-to-maturity (YTM) from the volume-weighted transaction price using the QuantLib library. We compute bond credit spreads as the bond yield-to-maturity minus the matched-maturity Treasury yield from the Liu and Wu (2021) zero-coupon yield curve, interpolated across seven key-rate tenors.[3] We drop missing observations, winsorize observations at the 1^{st} and 99^{th} percentiles, and aggregate daily data to monthly volume-weighted credit spreads.

Credit spread data are merged with bond characteristics from Mergent FISD and firm-level data from Compustat and CRSP via six-digit base CUSIP, excluding all financial firms (SIC codes 6000–6999). The final sample consists of 308,248 bond-month observations covering 10,245 bonds issued by 800 firms across 12 Fama and French (1997) industries. The pre-ChatGPT period accounts for

[3] Liu and Wu (2021) employs a non-parametric kernel-smoothing approach to construct the term structure, which produces substantially lower pricing errors and better captures local 'kinks' as compared to other approaches in the literature (Gürkaynak, Sack, and Wright, 2007).

152,626 observations (49.5% of the sample), and the post-release for the remaining 155,622.

Panel A of Table 1 reports summary statistics for all variations of credit spreads used in our tests. The mean spread of 149 basis points and standard deviation of 113 basis points reflect the right-skewed distribution typical of samples spanning both investment-grade and high-yield bonds (Campbell and Taksler, 2003; Gilchrist and Zakrajšek, 2012). The unconditional mean spread declines from approximately 162 basis points in the pre-period to 135 basis points in the post-period, reflecting broader credit market conditions. Importantly, this aggregate decline masks the cross-sectional heterogeneity across firms with varying AI exposure that is the focus of our analysis.

2.2 Generative AI exposure

Our key independent variable is firm-level workforce exposure to generative AI from Eisfeldt, Schubert, and Zhang (2026), constructed by mapping occupation-level AI exposure scores from Eloundou, Manning, Mishkin, and Rock (2024) to the firm level using employment-weighted averages from Revelio Labs. The measure captures the fraction of a firm's workforce engaged in tasks susceptible to generative AI substitution or augmentation, and is standardized to mean zero and unit standard deviation. Panel B reports summary statistics. The raw mean of 0.341, standard deviation of 0.055, and interquartile range of 0.061 are consistent with Eisfeldt, Schubert, and Zhang (2026), confirming our sample is comparable to theirs. Critically, we assign each firm its exposure score as of the pre-period and hold it fixed throughout, ensuring the measure captures structural exposure at the time of ChatGPT's release rather than firms' subsequent response to it.

2.3 Firm and bond characteristics

Panel C of Table 1 reports summary statistics for firm-level variables, following the credit spread determinants literature (Collin-Dufresne, Goldstein, and Martin, 2001; Campbell and Taksler, 2003; Gilchrist and Zakrajšek, 2012). The median firm has total assets of nearly \$59 billion and leverage of approximately 40%, reflecting the large-capitalization, moderately-levered profile typical of TRACE-eligible issuers. Tangibility averages 40% with a standard deviation of 27%, consistent with our sample spanning diverse industries (listed in Table A2 in Appendix A1). Panel D confirms that the average bond carries a coupon of 4.45%, a remaining maturity of 11.4 years, and an S&P rating of approximately BBB+. Summary statistics for all remaining variables are consistent with comparable samples in similar studies (Campbell and Taksler, 2003; Gilchrist and Zakrajšek, 2012; Faulkender and Petersen, 2006).

Having described the data and variable construction, we next examine whether workforce exposure to generative AI is reflected in corporate credit spreads.

3 Baseline results

We begin our empirical analysis with event-study and portfolio-sort evidence establishing that high- and low-exposure firms followed parallel credit spread trajectories prior to ChatGPT's release. We then estimate the baseline difference-in-differences specification and assess robustness to alternative fixed-effect structures, dependent variable definitions, and sample restrictions.

3.1 Identification

We estimate the effect of generative AI exposure on corporate bond credit spreads using a difference-in-differences design that exploits cross-sectional variation in firms' workforce exposure to generative AI. Our baseline specification is:

$$\text{Spread}_{i,b,t} = \beta \cdot \text{Post}_t \times \text{GenAI Exp}_i + \alpha_b + \gamma_t + \mathbf{X}'_{i,b,t}\delta + \varepsilon_{i,b,t} \tag{1}$$

where $\text{Spread}_{i,b,t}$ is the credit spread in basis points on bond b issued by firm i in month t, measured as the yield-to-maturity minus the matched-maturity Treasury yield interpolated from the Liu and Wu (2021) yield curve. Post_t is an indicator equal to one for year $\geq$ 2023, capturing the period following ChatGPT's public release in November 2022. GenAI Exp_i is the firm-level workforce exposure to generative AI from Eisfeldt, Schubert, and Zhang (2026), standardized to mean zero and unit standard deviation. The coefficient of interest is β, which measures the differential change in credit spreads for high-exposure firms following ChatGPT's release, relative to low-exposure firms. α_b are bond fixed effects, γ_t are year-month fixed effects, and $\mathbf{X}_{i,b,t}$ is a vector of time-varying firm and bond controls including assets, ROA, leverage, Tobin's Q, tangibility, liquidity, return, cash flow volatility, remaining maturity, and trading volume. All variables are as defined in the data section. In all regressions, standard errors are clustered at the firm level.

We begin by checking if the parallel trends assumption holds. Figure 1 plots the monthly event-study coefficients $\hat{\beta}_\tau$ from a version of equation (1) in which the single Post_t dummy is replaced by a set of 24 month-specific dummies from 12 months before to 12 months after ChatGPT's release. The dummy for month -3 is the omitted category. Each coefficient measures the differential credit spread between high- and low-exposure firms at horizon τ relative to November 2022.

The pre-treatment coefficients spanning months -12 through -4 are uniformly close to zero and are in almost all cases not statistically significant at conventional levels. High- and low-exposure firms followed parallel credit spread trajectories in the twelve months preceding ChatGPT's release, supporting the identifying assumption of the difference-in-differences design. The only anomalous coefficient is the one for month -5, estimated to be -3.3 basis points with a statistic of -2.60. From the figure this appears to be an isolated fluctuation with no systematic trend in any of the remaining pre-months. The negative significant coefficient in month -5, corresponding to June 2022, may reflect the Federal Reserve's 75 basis points increase in the Federal Funds rate, which Gilchrist and Zakrajšek (2012) and others document as a driver of credit spread widening.

Figure 1 also shows that the post-treatment pattern is markedly different. The coefficient turns positive and significant beginning at month 0, rising from 3.4 basis points at ChatGPT's release to a sustained level of 3.5 to 5.6 basis points through the end of the sample period. The increase in spreads for high-exposure firms is not transient and persists throughout the twelve-month post-period, consistent with a persistent repricing of AI disruption risk rather than a short-lived market reaction to a single news event. The effect even intensifies modestly at longer horizons, reaching approximately 6 basis points by month eight, suggesting that credit markets continued to incorporate new information about generative AI adoption risk as more information accumulated over the sample period.

Figure 2 complements the event-study evidence in Figure 1 by examining the time-series behavior of credit spreads of bond portfolios sorted by GenAI exposure. At the start of each year, we sort firms into quintiles based on their GenAI exposure score and compute the value-weighted average credit spread using data for all bonds in a quintile for the lowest-exposure quintile (Q1) and the highest-exposure quintile (Q5). To abstract from aggregate spread movements — which declined substantially over the sample period as monetary policy normalized — we demean each quintile's spread series by its pre-ChatGPT mean, computed over January 2021 through October 2022. The figure therefore plots the change in credit spreads relative to each quintile's own pre-period baseline.

Prior to ChatGPT's release in November 2022, both quintiles track closely to each other with no systematic divergence between high- and low-exposure firms. This provides visual confirmation of the parallel pre-trends, consistent with the event-study evidence in Figure 1 which also shows parallel trends. Following ChatGPT's release, the two quintiles diverge sharply. The high-exposure quintile (Q5) rises above its pre-period baseline while the low-exposure quintile (Q1) continues to decline, consistent with a broad tightening of aggregate spreads that does not seem to benefit high-exposure firms. By the end of our sample period, Q5 spreads are approximately 13.6 basis points above their pre-ChatGPT mean while Q1 spreads are approximately 27.1 basis points below theirs — a total difference of approximately 41 basis points that is consistent with our hypothesis regarding the persistent repricing of AI disruption risk for high-exposure issuers. In unreported results, we confirm that this pattern is robust to alternative schemes for computing portfolio spreads such as equal-weighting the portfolios.

Taken together, the event-study and portfolio-sort evidence establish that the post-ChatGPT widening of credit spreads for high-exposure firms is not an artifact of aggregate conditions, pre-existing trends, or bond-level heterogeneity. Next, we proceed to quantify and confirm this effect formally using a difference-in-difference framework.

3.2 Baseline results

Table 2 reports the baseline difference-in-differences estimates of equation (1). Each column in the Table presents the results for a different specification. Thus, column (1) includes bond fixed effects only; column (2) adds year-month fixed effects, which absorb all macroeconomic variation common to bonds in the same calendar month, including Federal Reserve policy changes, aggregate credit market conditions, and economy-wide uncertainty shocks; column (3) replaces year-month fixed effects with year fixed effects; and column (4) replaces bond fixed effects with firm fixed effects. All regressions include the full set of control variables defined above.

Across all specifications, the coefficient on Post $\times$ GenAI Exp is positive, economically meaningful, and statistically significant at the 1% level or better. For instance, in column (2), a one-standard-deviation increase in GenAI exposure is associated with 6.2 basis points wider credit spreads following ChatGPT's release, relative to firms with lower exposure. The coefficient is relatively stable across all specifications, ranging from 6.2 to 7.9 basis points, indicating that the result is not sensitive to the choice of fixed effect structure or empirical specification. The baseline result survives inclusion of any and all fixed effects such as the year-month fixed effects in column (2), which absorb any aggregate spread movements in the post-ChatGPT period, confirming that the effect reflects cross-sectional variation in GenAI exposure rather than a common post-treatment shock to all bonds.

Note that in Table 2 the coefficients on the control variables have expected signs. For example, leverage is the strongest predictor of credit spreads, with a one-unit increase associated with approximately 190–208 basis points wider spreads. Conversely, monthly stock returns enter negatively, consistent with the fact that equity and debt markets co-move and firm-specific news identically impacts equity returns and bond prices. Although cash flow volatility is positive and significant in column (1), it loses any statistical significance once we control for year-month fixed effects, suggesting that this variable may be at least partially proxying for aggregate macroeconomic conditions.

The economic magnitude of the baseline result is meaningful. At 6.2 basis points per standard deviation of GenAI exposure, the effect exceeds the 4–6 basis point spread widening documented for policy uncertainty shocks by Kaviani, Kryzanowski, Maleki, and Savor (2020), consistent with generative AI representing a more pervasive source of firm-level uncertainty than episodic policy shocks. Scaled to the interquartile range of GenAI exposure — a difference of 1.033 standard deviations between firms at the 75^{th} and 25^{th} percentiles — the implied spread differential is approximately 6.4 basis points. Relative to the sample mean spread of 148.7 basis points, this represents a 4.3% increase in borrowing costs for high-exposure firms. Translated into dollar terms, the 8 basis point spread differential for a one-standard-deviation increase in exposure implies approximately \$1.39 billion in additional annual interest costs for the \$1.74 trillion in bonds outstanding for high-exposure firms in our sample — a magnitude with direct implications for corporate investment and capital structure decisions.

Table A3 in the Appendix decomposes total GenAI exposure into core and supplemental components following Eisfeldt, Schubert, and Zhang (2026). Core exposure captures tasks central to an occupation that generative AI can directly substitute; supplemental exposure captures peripheral tasks where generative AI complements rather than replaces workers. Core exposure alone generates a 5.5 basis point spread widening, statistically significant at the 1% level with a $t = 2.95$ (column 1). Supplemental exposure alone is not statistically significant ($t = 1.24$, column 2). When both are included jointly (column 3), core exposure remains significant at nearly 5.7 basis points, i.e., its magnitude is higher than in column 1, while supplemental exposure remains statistically insignificant. The baseline result is therefore driven by bondholders pricing the uncertainty surrounding whether high-exposure firms successfully navigate the transition as core job functions are replaced by AI.

Table 3 examines whether the baseline result extends beyond ChatGPT's release to the other instances of generative AI model releases or updates. Panel A reports stacked difference-in-differences estimates following Baker, Larcker, and Wang (2022) across five major GenAI releases: ChatGPT (November 2022), GPT-4 (March 2023), Llama 2 (July 2023), Gemini (December 2023), and GPT-4o (May 2024). By constructing a separate clean control group for each event cohort, the stacked design avoids the negative weighting problem that arises in standard two-way fixed effects when treatment effects are heterogeneous across cohorts (Goodman-Bacon, 2021). All regressions also include bond, firm, year-month, or year fixed effects interacted with a cohort fixed effect. The cohort fixed effect is defined using data from the period three months before to three months after each event. The coefficient on Post $\times$ GenAI Exp is positive and statistically significant at the 5% level or better across all specifications, ranging from 1.1 to 2.4 basis points. The smaller magnitude relative to Table 2 likely reflects averaging across five events, each with potentially different effect sizes. The results in this panel confirm that the credit spread widening is not driven by the ChatGPT event alone but also by other instances of generative AI model releases or updates.

Panel B presents results for a complementary test using a 'dose-response' specification in which the binary Post indicator is replaced by the cumulative count of major generative AI model releases or updates in each month, interacted with firm-level exposure. The cumulative count starts from zero before November 2022 to one after ChatGPT's release, two after GPT-4, and so on up to five after GPT-4o. The coefficient measures the average incremental effect on credit spreads associated with each additional instance of generative AI model release or update per standard deviation of firm-level exposure, under the assumption that at the margin, each event contributes equally the documented effect.

The coefficient in Panel B ranges from 1.6 to 2.7 basis points across various specifications, all statistically significant at the 1% level throughout. To interpret these coefficients, consider two firms: firm H at the 75^{th} percentile of standardized GenAI exposure (+0.291 SD) and firm L at the 25^{th} percentile (-0.742 SD), a difference of 1.033 SD. Before any GenAI event (CumEvents = 0), the model predicts no differences in credit spreads between them. After ChatGPT's release (CumEvents = 1), the predicted spread differential between H and L is $2.708 \times 1.1 \approx 3.0$ basis points. After GPT-4 (CumEvents = 2), it widens to $2.708 \times 2 \times 1.1 \approx 6.0$ basis points, and so on. The coefficient therefore captures the average per-event increment in the spread differential between high- and low-exposure firms, under the assumption that each event contributes equally at the margin. The positive and significant coefficient on the cumulative event interaction is consistent with credit markets incorporating information about AI disruption risk progressively across successive GenAI milestones, rather than reacting exclusively to ChatGPT's release as a one-time shock. We next assess its robustness to alternative sample definitions, dependent variable constructions, and empirical specifications.

3.3 Robustness

Table 4 examines whether our baseline result is robust to alternative empirical specifications. Panel A addresses the concern that the effect reflects a short-lived announcement reaction after ChatGPT's release rather than a persistent repricing of credit risk. Dropping the two months immediately surrounding ChatGPT's release (November–December 2022) leaves the coefficient essentially unchanged at 6.6 basis points ($t = 3.05$). Dropping the November 2022 through January 2023 window produces an identical result (6.6 basis points, $t = 3.00$). Thus, the baseline finding is not driven by a transient announcement spike.

Panel B replaces the Treasury-based credit spread with two alternative measures of credit spreads. Using the Gilchrist and Zakrajšek (2012) credit spread — constructed against a duration-matched synthetic Treasury bond —- delivers a coefficient of 5.0 basis points ($t = 2.73$). The estimate is higher at 5.7 basis points ($t = 2.19$) when using the spread over AAA-rated non-financial corporate bond yields as the risk-free benchmark. Both confirm that the result is not an artifact of benchmarking bond yields against Treasury yields computed from Liu and Wu (2021) or of flight-to-quality effects that may inflate Treasury-based spreads in the post-ChatGPT period.

In Panel C, we aggregate bonds to the firm-month level and augment the specification with increasingly demanding industry-specific time fixed effects. The baseline firm-month specification produces 10.93 basis points ($t = 4.41$), consistent with the bond-level result after removing within-firm bond heterogeneity. Adding Fama and French (1997) 49-industry $\times$ year fixed effects, which absorb any annual industry-level trends in credit spreads, reduces the coefficient to nearly 7.40 basis points ($t = 2.34$). Replacing these with FF49-industry $\times$ year-month fixed effects — the most demanding specification, absorbing monthly industry-level variation — leaves the coefficient at 7.24 basis points ($t = 2.24$). This robustness test is particularly relevant given that the Eisfeldt, Schubert, and Zhang (2026) exposure measure varies substantially across industries by construction. The stability of the coefficient under industry $\times$ time fixed effects confirms the effect is identified from within-industry variation in firm-level exposure, ruling out the possibility that technology-adjacent industries simply faced wider spreads post-ChatGPT for reasons unrelated to generative AI.

Finally, Panel D restricts the sample to non-callable bonds, which carry no embedded call option that could mechanically dampen spread sensitivity to credit risk changes. The coefficient rises sharply to approximately 20 basis points across all four fixed effect structures (t ranging from 3.77 to 3.90), roughly three times the full-sample baseline result in Table 2. The amplification reflects the fact that callable bonds allow issuers to refinance their debt when spreads tighten, introducing an optionality component that compresses observed spread movements and biases the full-sample estimate toward zero. The non-callable subsample, although considerably smaller, reveals the underlying credit risk effect in the absence of this.

These baseline results establish that high-exposure firms face relatively wider credit spreads following ChatGPT's release, a pattern that is robust across specifications and consistent with our identification assumptions. We now develop a theoretical framework that rationalizes this finding and generates additional cross-sectional predictions that we take to the data in Section 5.

Figure 3 reports the results of Fisher randomization inference. We randomly reassign GenAI exposure scores across firms 1,000 times, re-estimating the preferred specification after each reassignment, and plot the resulting permutation distribution of coefficients. Under the null hypothesis that GenAI exposure has no effect on credit spreads, randomly shuffling exposure across firms should produce coefficients of similar magnitude to the actual estimate. The permutation distribution is centered tightly around zero with a standard deviation of 1.75 basis points, reflecting the absence of any spurious correlation between the cross-sectional structure of the data and credit spread dynamics. The actual coefficient of 6.15 basis points exceeds all 1,000 permuted coefficients, placing it at the 100^{th} percentile of the distribution ($p < 0.001$). The randomization evidence conclusively rejects the null hypothesis and confirms that the baseline result cannot be attributed to chance or to any mechanical feature of the data generating process.[4]

4 Theoretical framework

In this section, we develop a labor-technology substitution model with parameter uncertainty. The model extends Zhang (2019) and Knesl (2023) by introducing an ingredient absent from those frameworks: at the moment a new technology arrives, its fundamental adoption parameters are unknown to firms and investors. We embed this parameter uncertainty in the structural credit risk framework of Merton (1974) to derive the joint behavior of equity values and credit spreads. The model nests Zhang (2019) as a special case when parameter uncertainty is absent, and shows that it is precisely this uncertainty, and not the substitution option itself, that generates the equity-debt divergence: high-exposure firms experience relatively wider spreads than low-exposure firms, while simultaneously capturing equity gains. We describe the model below, with formal derivations provided in the Appendix A2.

4.1 Setup

Consider a firm in which a fraction $\lambda \in [0, 1]$ of the workforce performs tasks potentially exposed to generative AI as measured by Eisfeldt, Schubert, and Zhang (2026). The firm is financed by both equity and debt. Debt takes the form of a single zero-coupon bond with face value F and maturity $\mathcal{T}$. The value of the firm today (at $t = 0$) equals the present value of all future cash flows. Under the structural framework of Merton (1974), this implies that equity can be valued as a European call option on firm assets and debt as a portfolio of a risk-free bond and a short

[4] In a complementary test, we conduct a pre-period placebo test by assigning a placebo treatment date one year prior to ChatGPT's release (November 2021) and re-estimating the preferred specification. The resulting coefficient is near-zero and statistically insignificant, confirming that the result is not driven by a pre-existing differential trend in credit spreads between high- and low-exposure firms that predates ChatGPT's release.

European put:

$$E_0 = V_0 \cdot N(d_1) - F \cdot e^{-r\mathcal{T}} \cdot N(d_2) \tag{2}$$

$$D_0 = F \cdot e^{-r\mathcal{T}} - [F \cdot e^{-r\mathcal{T}} \cdot N(-d_2) - V_0 \cdot N(-d_1)] \tag{3}$$

where $N(\cdot)$ is the standard normal CDF, r is the risk-free rate, and d_1 and d_2 are defined as:

$$d_1 = \frac{\ln(V_0/F) + (r + \sigma^2/2)\mathcal{T}}{\sigma\sqrt{\mathcal{T}}}, \qquad d_2 = d_1 - \sigma\sqrt{\mathcal{T}} \tag{4}$$

As a result, the credit spread is given by:

$$CS = -\frac{1}{\mathcal{T}} \cdot \ln\left[N(d_2) + \frac{V_0}{F \cdot e^{-r\mathcal{T}}} \cdot N(-d_1)\right] \tag{5}$$

Prior to the arrival of the technological shock, firm asset value follows a geometric Brownian motion $dV/V = \mu_0\, dt + \sigma_0\, dW_t$ under the physical measure $\mathbb{P}$, where μ_0 is the physical drift, σ_0 is the baseline asset volatility, and W_t is a standard Brownian motion. Under the risk-neutral measure $\mathbb{Q}$ used for derivative pricing, the drift is replaced by r and both equity and debt are priced as functions of the volatility σ. Equity, as the long call, benefits from higher σ, while bondholders, holding the short put are harmed. Next, we introduce the adoption option created by ChatGPT's release and the parameter uncertainty that distinguishes our setting from the existing literature.

4.2 The adoption option and parameter uncertainty

We assume that the generative AI shock represented by ChatGPT's release arrives at time $t = 0$ and creates a real option to adopt the technology and restructure its AI-exposed workforce. Following Zhang (2019), the option delivers a perpetual flow payoff of $\delta \cdot w \cdot \lambda$ per unit time if successfully exercised, where $\delta > 0$ is the expected productivity gain per unit of AI-exposed labor and w is the wage rate.[5] If the firm adopts the technology, it must pay a one-time fixed cost of

[5] We model the adoption payoff as a permanent flow saving discounted as a perpetuity, so that $p \cdot (\delta \cdot w \cdot \lambda/r)$ is the present value of expected cost savings under a constant discount rate r. This is a standard simplification in the

$\kappa > 0$. Adoption succeeds with probability $p \in (0, 1)$. The net present value of adoption under $\mathbb{P}$ is given by:

$$\Phi = p \cdot \left(\frac{\delta \cdot w \cdot \lambda}{r} \right) - \kappa > 0 \tag{6}$$

For concreteness, we assume $\Phi > 0$, as that is the empirically relevant case, consistent with the positive equity market reaction to the arrival of generative AI technology documented by Eisfeldt, Schubert, and Zhang (2026). Our model's predictions for credit spreads hold for any value of Φ.

Given this setup, the standard optimal adoption rule under certainty applies as in Dixit and Pindyck (1994). With irreversible fixed cost κ, the firm faces a classic real options problem, and chooses to wait until the value of immediate adoption Φ exceeds the continuation value of keeping the option alive, generating a unique adoption threshold $\Phi^* > 0$. The firm adopts immediately if $\Phi \geq \Phi^*$ and waits otherwise. The adoption threshold Φ^* is a closed-form function of κ, p, δ, r, and the volatility of the underlying cash flow process (Dixit and Pindyck, 1994).[6]

Crucially, computing Φ^* requires knowing the values of κ, p, and δ, or at a minimum forming precise beliefs about them. When parameters are known with certainty, this is of course the case. We can then compute the exercise threshold and implement the optimal adoption rule. That is, the firm can clearly determine when $\Phi \geq \Phi^*$ and chooses to adopt the technology. In this case, the adoption decision is well-defined with a determinate outcome.

The critical departure we make from the standard setup from Zhang (2019) and Knesl (2023) is that at time $t = 0$, the parameters governing the adoption decision (κ, p, δ) are not known with certainty. Firms and investors instead observe noisy signals:

$$\hat{\kappa} = \kappa + \varepsilon_\kappa, \qquad \hat{p} = p + \varepsilon_p, \qquad \hat{\delta} = \delta + \varepsilon_\delta \tag{7}$$

where $\varepsilon_\kappa, \varepsilon_p, \varepsilon_\delta$ are noise terms with mean zero and variances $\sigma_\kappa^2, \sigma_p^2, \sigma_\delta^2$, respectively. We let

real options literature (Dixit and Pindyck, 1994) and does not affect the credit spread predictions, which depend exclusively on asset volatility σ under $\mathbb{Q}$ rather than on the drift and level of cash flows.

[6] Φ^* does not depend on the wage rate as we treat this as fixed. In a richer model where wage rates are stochastic, w would enter the dynamics of Φ and therefore affect Φ^*. Since our model is static at $t = 0$ and we are interested in the cross-sectional implications of λ rather than wage dynamics, this simplification does not affect any of our results.

the vector $\Sigma = (\sigma_\kappa, \sigma_p, \sigma_\delta)$ characterize total parameter uncertainty, with $\|\Sigma\| = \sqrt{\sigma_\kappa^2 + \sigma_p^2 + \sigma_\delta^2}$. Thus, $\|\Sigma\|$ aggregates the three uncertainty sources under the simplifying assumption that ε_κ, ε_p, and ε_δ are mutually independent, an assumption made for tractability.[7]

Equation (7) simply captures the intuition that because generative AI had no established adoption track record at $t = 0$, the true values of adoption cost, success probability, and productivity gain were genuinely unknown to firms and investors alike. This prevents implementation of the first-best adoption rule: with only noisy signals $(\hat{\kappa}, \hat{p}, \hat{\delta})$, firms cannot evaluate whether $\Phi \geq \Phi^*$ holds, making the optimal adoption decision indeterminate. Facing these noisy signals, investors form prior beliefs over the true parameter values (κ, p, δ) and update them continuously as new information about the technology's viability arrives (Veronesi, 1999). Next, we derive the asset pricing implications of this belief revision process — specifically, how continuous belief updating injects an additional diffusion term into the asset value process under $\mathbb{Q}$, increasing effective risk-neutral volatility proportionally to workforce exposure λ.

4.3 Asset value dynamics with parameter uncertainty

With parameter uncertainty $\|\Sigma\| > 0$, two distinct effects arise for the firm. First, the firm's expected cash flows rise under $\mathbb{P}$ because successful adoption raises the firm's physical drift $\Delta\mu > 0$. This increases total firm value and benefits equity as the residual claim on firm assets. Second, because the adoption parameters (κ, p, δ) are unknown, uncertainty propagates from the parameters into the payoff itself: the realized adoption payoff $\hat{\Phi}$ is not a point estimate but a random variable whose variance is strictly increasing in both $\|\Sigma\|$ and λ. We elaborate on each effect below.

Impact on cash flows. Under $\mathbb{P}$, expected cash flows increase because successful adoption (probability p) generates savings of $\delta \cdot w \cdot \lambda$ per unit time, raising the expected drift to $\mu_0 + \Delta\mu$, where $\Delta\mu = p \cdot \delta \cdot w \cdot \lambda / V_0 > 0$. Higher expected cash flows raise total firm value and therefore equity

[7] A more general treatment would replace $\|\Sigma\|^2$ with a quadratic form $\Sigma'\Omega\Sigma$, where Ω represents the full covariance matrix of parameter uncertainty, allowing option parameters to covary. For instance, uncertainty about productivity could covary positively with uncertainty about the adoption cost if both variables reflect the ignorance about the maturity of the underlying technology. This generalization, even though more accurate, leaves the qualitative predictions from our model unchanged. Our key qualitative results hold under any positive definite Ω.

value, consistent with the positive equity returns documented by Eisfeldt, Schubert, and Zhang (2026). This effect operates under $\mathbb{P}$ and does not affect the pricing of either claim under $\mathbb{Q}$, where the drift is replaced by r.

Impact on volatility. Because adoption parameters are unknown, investors face a signal extraction problem. Under certainty, the adoption payoff is a single known number $\Phi = p \cdot (\delta \cdot w \cdot \lambda / r) - \kappa$ — a point estimate with no distribution. Conversely, when parameters are unknown, each of (κ, p, δ) is replaced by its noisy signal and the realized payoff from option exercise becomes:

$$\hat{\Phi} = \hat{p} \cdot \left(\frac{\hat{\delta} \cdot w \cdot \lambda}{r} \right) - \hat{\kappa} \tag{8}$$

where $\hat{p} = p + \varepsilon_p$, $\hat{\delta} = \delta + \varepsilon_\delta$, and $\hat{\kappa} = \kappa + \varepsilon_\kappa$. Because each parameter is uncertain, $\hat{\Phi}$ is a random variable spread around the expected value Φ. By a first-order Taylor expansion, its variance is:

$$\mathrm{Var}(\hat{\Phi}) \approx \left(\frac{\delta w \lambda}{r} \right)^2 \sigma_p^2 + \left(\frac{p w \lambda}{r} \right)^2 \sigma_\delta^2 + \sigma_\kappa^2 \tag{9}$$

That is, $\hat{\Phi}$ inherits the variance of all three noisy parameters and uncertainty about the parameters propagates into uncertainty about the real option payoff itself. $\mathrm{Var}(\hat{\Phi})$ is strictly increasing in each component of $\|\Sigma\|$ and in λ. Some realizations of $\hat{\Phi}$ will be much better than Φ, and others much worse; the distribution fans out around the expected value as $\|\Sigma\|$ and λ grow. Firms with greater AI workforce exposure face a larger adoption option and therefore larger payoff variance — the uncertainty scales with λ. Investors must now price a distribution of adoption outcomes rather than just a point estimate, injecting additional volatility into the asset value process under $\mathbb{Q}$. Following Veronesi (1999), as investors continuously update beliefs about the technology upon observing noisy signals, this belief revision injects an additional diffusion term into the asset value process under $\mathbb{Q}$, increasing effective risk-neutral volatility:

$$\sigma(\lambda, \Sigma) = \sigma_0 + \lambda \cdot \|\Sigma\| \tag{10}$$

The term $\lambda\|\Sigma\|$ is the incremental volatility attributable to parameter uncertainty: it is zero when parameters are known ($\|\Sigma\|= 0$), and strictly increasing in both the degree of parameter uncertainty $\|\Sigma\|$ and the firm's AI workforce exposure λ. Under $\mathbb{Q}$, both equity and debt are priced as derivative claims on firm assets. Equity, as a long call, benefits from higher σ ($\partial E/\partial\sigma > 0$): equity value rises. Debt, as a short put, is harmed by higher σ ($\partial D/\partial\sigma < 0$): bond value falls and credit spreads widen. The cash flow effect raises equity only; the volatility effect moves both claims in opposite directions, generating the equity-debt divergence.

Nesting Zhang (2019) as a special case. Our model nests Zhang (2019) as a special case. In Zhang (2019), firms face a technology adoption option with known parameters ($\|\Sigma\|= 0$) that arrives countercyclically and hedges macroeconomic risk and reduces the effective discount rate on bond cash flows. Known parameters mean the first-best adoption rule is implementable in the sense that firms can observe (κ, p, δ) directly, compute the optimal adoption threshold Φ^*, and determine with certainty whether $\Phi \geq \Phi^*$ holds — the adoption decision is therefore determinate and generates no additional uncertainty about future asset values. Consequently, no additional diffusion term is injected into the asset value process under $\mathbb{Q}$:

$$\|\Sigma\|= 0 \implies \Delta\sigma^{\mathbb{Q}} = 0 \tag{11}$$

Combined with the countercyclical hedge — i.e., the fact that the option pays off precisely when aggregate conditions are bad and marginal utility is high, reducing firm exposure to systematic risk and the risk premium bondholders require — credit spreads tighten and the substitution channel dominates. In our setting, ChatGPT's release satisfies $\|\Sigma\|\gg 0$ and generative AI's payoff is assumed to be uncorrelated with aggregate economic conditions. In other words, the option arrives unconditionally rather than countercyclically, rendering the countercyclical hedge inoperative. Parameter uncertainty $\|\Sigma\|> 0$ is therefore both necessary and sufficient for the equity-debt divergence and is precisely the ingredient that converts a spread-tightening substitution option into a spread-widening disruption shock.

Zhang (2019) does not examine credit spreads directly, but the implication of the framework is unambiguous — known-parameter options leave risk-neutral volatility unchanged and therefore leave spreads unaffected: without parameter uncertainty, any real option raises equity but leaves credit spreads unchanged. With it, the belief revision process injects additional volatility under $\mathbb{Q}$, simultaneously raising equity value and generating relatively wider credit spreads for high-exposure firms. Setting $\|\Sigma\|= 0$ recovers Zhang (2019)'s result; setting $\|\Sigma\|> 0$ generates ours.

Disruption versus substitution. The cash flow and volatility dynamics described in this section map directly onto the channels introduced in Section . The total cross-sectional spread differential decomposes as:

$$dCS = \underbrace{\frac{\partial CS}{\partial \mu} \cdot \Delta\mu}_{\text{Substitution=0}} + \underbrace{\frac{\partial CS}{\partial \sigma^{\mathbb{Q}}} \cdot \Delta\sigma^{\mathbb{Q}}}_{\text{Disruption>0}} \tag{12}$$

Under $\mathbb{Q}$, the drift μ is replaced by r, so the Substitution term is zero and credit spreads respond exclusively to the volatility effect $\Delta\sigma^{\mathbb{Q}} = \lambda\|\Sigma\|$. The disruption channel dominates whenever $\|\Sigma\|> 0$; setting $\|\Sigma\|= 0$ eliminates the disruption channel entirely. Under the Zhang (2019) framework, spreads may tighten through the countercyclical hedge mechanism, which falls outside the scope of this paper.

4.4 A role for governance

In addition to modifying the basic framework in Zhang (2019) to account for parameter uncertainty, we extend the baseline model described above to incorporate a role for firm corporate governance quality $g \in [0, 1]$, where higher values of g reflect mechanisms such as board oversight, CEO incentive alignment, and institutional monitoring that can credibly constrain the tendency of equity-aligned management to make opportunistic adoption decisions.

Beyond constraining opportunistic adoption decisions, stronger governance also mitigates the agency costs associated with the free cash flow released by successful AI adoption: firms that successfully adopt will generate cost savings that, absent effective monitoring, may be deployed in value-destroying investments or extracted as private benefits rather than returned to investors (Jensen, 1986). Higher g therefore dampens the spread-widening effect through two complementary mechanisms — reducing parameter uncertainty about the adoption decision itself and disciplining the deployment of the cash flows that adoption generates. Both lead to a narrowing of the set of outcomes bondholders must price, reducing effective parameter uncertainty:

$$\sigma(\lambda, \Sigma, g) = \sigma_0 + \lambda \cdot \|\Sigma\| \cdot (1 - g) \tag{13}$$

Higher g reduces effective volatility and dampens the spread-widening effect. In our empirical analysis, we use CEO portfolio delta and institutional ownership stability as empirical proxies for g. With the model specification complete, we next describe the empirical implications of our framework.

4.5 Empirical predictions from the model

We now state the propositions from our model that map directly onto our empirical tests. All predictions follow from the model setup with formal proofs in Appendix A2.

Proposition 1 (Equity-debt divergence)**.** *For any firm with $\lambda > 0$ and $\|\Sigma\| > 0$, ChatGPT's release simultaneously raises equity value through both the cash flow effect ($\partial E/\partial\mu > 0$ under $\mathbb{P}$) and the volatility effect ($\partial E/\partial\sigma > 0$ under $\mathbb{Q}$), while generating relatively wider credit spreads compared to low-exposure firms through the volatility effect alone ($\partial CS/\partial\sigma > 0$ under $\mathbb{Q}$). Total firm value rises but is redistributed toward equity.*

Discussion. The equity-debt divergence is our central contribution. Under $\mathbb{P}$, equity benefits from both higher expected cash flows and the option value of uncertain adoption. Under $\mathbb{Q}$, only the volatility increase affects spreads — shareholders price the benefits of the adoption option while bondholders price the parameter uncertainty surrounding its exercise. The spread differential is increasing in λ: firms with greater AI workforce exposure hold a bigger adoption option, which generates a greater incremental rise in asset volatility. A firm with no workforce exposure ($\lambda = 0$) experiences no incremental increase in volatility and serves as the implicit control group. Setting $\|\Sigma\| = 0$ eliminates the equity-debt divergence entirely.

Proposition 2 (Implied equity volatility)**.** *ChatGPT's release increases implied equity volatility for high-exposure firms relative to low-exposure firms, reflecting the immediate incorporation of parameter uncertainty* $\|\Sigma\|$ *into option prices under* $\mathbb{Q}$.

Discussion. We test whether the introduction of the adoption option with parameter uncertainty increases implied equity volatility. Implied equity volatility is the right empirical proxy as it is forward-looking by construction and immediately incorporates the full distribution of adoption outcomes along with investor uncertainty about (κ, p, δ). The model predicts a cross-sectional increase: firms with high workforce exposure will face a larger parameter uncertainty shock $\lambda\|\Sigma\|$ and therefore a larger increase in implied volatility relative to low-λ firms. This is a direct market-based test of the risk-neutral volatility increase that drives the credit spread result. Analyst forecast dispersion provides a complementary measure of forward-looking uncertainty: when adoption parameters are unknown, even informed analysts disagree about future earnings, so dispersion should rise for high-exposure firms alongside implied volatility.

Proposition 3 (Left tail asymmetry)**.** *Parameter uncertainty generates an asymmetric increase in risk-neutral volatility that is concentrated in the left tail of the asset value distribution: implied volatility on deep out-of-the-money put options increases for high-exposure firms relative to low-exposure firms following ChatGPT's release.*

Discussion. The upside from adoption is bounded by $\delta \cdot w \cdot \lambda/r$ and the downside from failed adoption combined with competitive displacement is unbounded. This asymmetry implies that parameter uncertainty $\|\Sigma\|$ fattens the left tail of the asset value distribution disproportionately, increasing the risk-neutral probability of large negative outcomes that deep OTM puts protect against. High-exposure firms therefore exhibit a larger increase in deep OTM put implied volatility than low-exposure firms following ChatGPT's release.

Proposition 4 (Amplification under distress)**.** $\partial^2(CS)/\partial\sigma\,\partial L > 0$, *where* $L = Fe^{-r\mathfrak{T}}/V_0$. *The sensitivity of the credit spread to the parameter uncertainty shock is increasing in leverage because the Merton credit spread is convex in volatility near the default boundary.*

Discussion. Firms with high leverage or greater credit fragility sit close to the default boundary; when $\|\Sigma\|$ widens the asset value distribution, these firms face disproportionately greater left-tail or downside risk and may have less financial slack to absorb failed adoption or competitive displacement. In our empirical tests, we use the firm's S&P rating, Altman Z-score, equity volatility, and leverage as proxies for credit fragility.

Proposition 5 (Impact of governance)**.** $\partial(CS)/\partial g = [\partial(CS)/\partial\sigma] \cdot \lambda\|\Sigma\|\cdot(-1) < 0$. *Higher g reduces effective parameter uncertainty $\|\Sigma\|(1-g)$, moderating the risk-neutral volatility increase.*

Discussion. CEO equity incentives (high delta) align management with long-run value maximization. Stable institutional ownership constrains value-extracting adoption behavior—for example, adopting AI primarily to justify layoffs rather than to maximize productivity. In our empirical tests, we proxy for governance quality using CEO delta and institutional ownership stability.

4.6 Calibration

Having established the model's theoretical predictions, we now calibrate its key parameters to assess whether the implied equity and credit spread effects are consistent with observed data. Table 5 reports the parameterization of the model.

Most of our model parameters are selected to match the summary statistics in Table 1. For our baseline case, we set asset volatility to $\sigma_0 = 0.293$ equal to the sample mean equity volatility, consistent with the Merton (1974) framework in which equity volatility proxies for asset volatility which itself is unobserved (Campbell and Taksler, 2003). The leverage ratio $L = 0.40$ also matches the sample mean reported in Panel C of Table 1, consistent with the moderately-levered profile of TRACE-eligible bond issuers (Collin-Dufresne, Goldstein, and Martin, 2001). Bond maturity $\mathfrak{T} = 7.4$ years is the sample median from Panel D. The interquartile range of raw GenAI workforce exposure $\Delta\lambda = 0.061$ is drawn from Panel B of Table 1 and represents the spread between a high-exposure firm at the 75^{th} percentile and a low-exposure firm at the 25^{th} percentile. Parameter uncertainty $\|\Sigma\|$ is the sole free parameter, calibrated to match the baseline empirical result as described below.

Credit spread prediction From Appendix A2, the cross-sectional spread differential between a firm with high and low workforce exposure to generative AI is:

$$\Delta CS = \frac{\partial(CS)}{\partial\sigma} \cdot \Delta\lambda \cdot \|\Sigma\| = \frac{n(d_2) \cdot \sqrt{\mathcal{T}} \cdot Fe^{-r\mathcal{T}}}{\mathcal{T} \cdot D_0} \cdot \Delta\lambda \cdot \|\Sigma\| \tag{14}$$

We set $\Delta CS = 8$ basis points per standard deviation of GenAI exposure, which matches our empirical estimate in Section 3 above and solve for the unobserved parameter uncertainty $\|\Sigma\|$. This yields an implied parameter uncertainty of $\|\Sigma\| = 0.120$, or approximately 12%: a 12% standard deviation in adoption parameter uncertainty reflects genuine but modest ignorance about the cost, success probability, and productivity consequences of a technology with no established adoption track record, which is economically plausible. For context, Kaviani, Kryzanowski, Maleki, and Savor (2020) report that a one-standard-deviation increase in the Baker, Bloom, and Davis (2016) Policy Uncertainty Index raises credit spreads by 9.6 basis points — broadly consistent with our 8 basis point estimate — suggesting that the parameter uncertainty implied by our model is of a similar order of magnitude to the uncertainty generated by major policy shocks. Across the interquartile range of GenAI exposure, the model predicts a spread differential of approximately 9 basis points.

Equity prediction The model-implied equity value change from the volatility effect alone per standard deviation of GenAI exposure is:

$$\frac{\Delta E}{E_0} = \frac{V_0 \cdot n(d_1) \cdot \sqrt{\mathcal{T}}}{E_0} \cdot \mathrm{SD}(\lambda) \cdot \|\Sigma\| \tag{15}$$

For the calibrated parameters, the volatility effect generates an equity gain of approximately 0.35% per standard deviation of exposure. This is directly comparable to the benchmark result from Eisfeldt, Schubert, and Zhang (2026), who report that a one-standard-deviation increase in GenAI exposure is associated with 1.7% higher cumulative abnormal returns over the ChatGPT event period (their Table V, Column 1). Our model-implied equity gain of 0.35% accounts for only 21% of this 1.7% total.

The difference between our model's prediction and the results in Eisfeldt, Schubert, and Zhang (2026) is due to the fact that the volatility effect driven by parameter uncertainty is not the primary driver of equity gains. Rather, it is the cash flow effect — higher expected earnings from successful AI adoption under $\mathbb{P}$ — that accounts for the remaining 79%, or approximately 1.56% per standard deviation of exposure. This interpretation is consistent with Eisfeldt, Schubert, and Zhang (2026)'s own evidence that cash flows drive their equity result: they document significant reductions in job postings for exposed occupations and increases in analyst forecasts of earnings per share following ChatGPT's release, confirming that the equity revaluation reflects expected future cost savings rather than a change in risk-neutral asset volatility.

The volatility effect contributes a small positive increment to equity while simultaneously generating the entire 8 basis point credit spread differential, since the cash flow effect does not enter credit spread pricing under $\mathbb{Q}$. The asymmetry between a large equity response and a modest spread differential is therefore a direct prediction of the model's $\mathbb{P}/\mathbb{Q}$ structure: the same parameter uncertainty $\|\Sigma\|= 0.120$ simultaneously rationalizes both the equity and credit spread results through different probability measures.

Figures 4 and 5 illustrate the model's comparative statics using the calibrated parameter values in Table 5. Figure 4 plots the credit spread differential ΔCS — the spread for a firm with workforce exposure to generative AI relative to a baseline firm for which we assume no workforce exposure. The differential is plotted as a function of parameter uncertainty, $\|\Sigma\|$, for three levels of generative AI workforce exposure λ, corresponding to the 25^{th}, 50^{th}, and 75^{th} percentiles of the Eisfeldt, Schubert, and Zhang (2026) exposure distribution.

With no parameter uncertainty ($\|\Sigma\|= 0$), the model nests Zhang (2019): the spread differential is zero regardless of exposure. As $\|\Sigma\|$ increases from zero, ΔCS rises with the rate of increase larger for higher-exposure firms. For any given level of parameter uncertainty, firms with greater generative AI workforce exposure face relatively wider credit spreads, consistent with Proposition 1.

In Figure 5, Panel A plots the spread differential between the median-exposure firm ($\lambda = 0.331$) and a zero-exposure firm. ΔCS rises with parameter uncertainty $\|\Sigma\|$, accelerating as parameter uncertainty increases. Panel B plots the relative equity gain $\Delta E/E_0$ for the same firms. In contrast to spreads, equity appreciation is present even with no parameter uncertainty ($\|\Sigma\|= 0$), driven entirely by the cash-flow effect: expected firm value rises under $\mathbb{P}$ because successful adoption raises the physical drift consistent with Eisfeldt, Schubert, and Zhang (2026). As $\|\Sigma\|$ increases, equity continues to appreciate modestly as higher effective volatility benefits the residual claimant through the standard vega effect. The divergence between the two panels — spreads rising sharply while equity rises gently — is the model's central prediction.

4.7 Summary

The comparative statics highlight the model's two key testable predictions. First, an option to adopt a transformative new technology lifts equity valuations through higher expected cash flows while the uncertainty surrounding the option generates relatively wider credit spreads. Parameter uncertainty is the precise ingredient that converts a labor-technology substitution option from a spread-tightening mechanism (Zhang, 2019) into a source of risk for bondholders that is amplified by credit fragility.

Second, both effects are amplified by the scale of the adoption option: workforce exposure λ and parameter uncertainty $\|\Sigma\|$ together determine the magnitude of the effect. Taken together, the same level of parameter uncertainty that generates a spread differential of approximately 8 basis points per standard deviation also rationalizes the equity results.

5 Testing model predictions

We now take the model's predictions to the data, testing each of the five propositions in turn using the same difference-in-differences framework established in Section 3.

5.1 Equity-bond divergence

Table 6 tests for the prediction that ChatGPT's release simultaneously raises equity values and widens credit spreads for high-exposure firms (Proposition 1). Each column in the table uses a distinct dependent variable, documenting the divergence and confirming the credit spread widening reflects genuine default risk repricing.

Column (1) shows that a one-standard-deviation increase in GenAI exposure is associated with 0.55 percentage points higher monthly stock returns following ChatGPT's release. Column (2) shows that Tobin's Q rises by nearly 0.11 for high-exposure firms. Both results are statistically significant at the 5% level or better, with $t = 4.21$ and $t = 2.03$, respectively. The results are also economically meaningful. The 0.55 percentage point monthly stock return differential compounds to approximately 6.6 percentage points on an annualized basis, comparable to the 1.7 percentage point cumulative abnormal return per standard deviation of exposure documented by Eisfeldt, Schubert, and Zhang (2026) over the two-week window around ChatGPT's release. Note that our monthly estimate over a longer horizon captures the persistent component of the equity revaluation rather than the initial announcement reaction. The Tobin's Q increase of 0.11 represents approximately 5% of the sample mean of 1.985, a meaningful increase in market-to-book valuation for firms whose median assets is nearly \$59 billion, implying an increase in market capitalization of roughly \$6 billion for the median firm in our sample. These two results establish that high-exposure firms are valued more highly by equity markets following ChatGPT's release, consistent with the cash flow effect under $\mathbb{P}$ raising expected firm value.

Columns (3)–(7) document the results for the debt side in Proposition 1. Since the baseline result in Table 2 already establishes the increase in credit spreads for the full sample of 308,248 bond-month observations, columns (3)–(7) restrict the sample to those firms for whom CDS data is available. Using CDS data allows us to decompose the total credit spread into the default risk and liquidity risk components following Longstaff, Mithal, and Neis (2005), who use the relatively liquid CDS spread to proxy for the default component of corporate bond yields and attribute the residual bond–CDS basis to bond-specific liquidity — a decomposition that requires matched CDS data and is therefore not feasible for the full sample.

Column (3) shows that for this group bond credit spreads widen by 10.7 basis points ($t = 3.32$). Columns (4) and (5) show that CDS spreads at the 1- and 3-year tenors widen significantly — by 7.4 ($t = 2.33$) and 8.9 ($t = 2.11$) basis points respectively — confirming the widening reflects an increase in perceived default risk over the 1 – 3 year horizon. Columns (6) and (7) report the bond–CDS basis and both coefficients are economically small and statistically insignificant at conventional levels, ruling out bond-specific liquidity as an alternative explanation for the increase in credit spreads following ChatGPT's release.[8]

Overall, the results in Table 6 confirm that the same firms, exposed to the same shock, experience rising equity valuations and rising credit spreads simultaneously — a divergence that follows directly from the asymmetric payoff structures of equity and debt claims in the presence of parameter uncertainty, and cannot be attributed to bond market illiquidity or measurement artifact.

5.2 Volatility channel and left tail risk

Proposition 2 predicts that parameter uncertainty raises implied equity volatility for high-exposure firms, while Proposition 3 predicts that this increase is concentrated in the left tail of the asset value distribution. Table 7 tests both predictions using option-implied measures from OptionMetrics.

Panel A tests the volatility channel. Column (1) shows that ATM implied volatility rises by 2.1 percentage points for high-exposure firms following ChatGPT's release. This result is statistically significant at the 1% level ($t = 5.12$), and is the strongest result in the table. Against a sample mean ATM implied volatility of 0.293, this coefficient represents a 7% increase and is economically meaningful. This result confirms that the option markets immediately incorporate the

[8]When using the 5-, 7-, and 10-year CDS as the dependent variable, the coefficients on the interaction term are positive but statistically insignificant (t-statistics of 1.59, 1.39, and 1.29 respectively), suggesting that credit markets view the disruption risk as concentrated in the near-to-medium term. This is consistent with the learning mechanism in Veronesi (1999) and Pástor and Veronesi (2009). As adoption costs, success probabilities, and competitive consequences of generative AI become better understood over longer horizons, parameter uncertainty $\|\Sigma\|$ dissipates, moderating the credit spread effect at maturities beyond three years.

wider distribution of adoption outcomes into their pricing, consistent with Proposition 2 precisely as Drechsler (2013) and Chen and Han (2023) predict: parameter uncertainty widens the risk-neutral distribution of asset values, and forward-looking option prices are the natural first venue in which this uncertainty is priced.

Column (2) shows that analyst forecast dispersion rises by 0.33 for high-exposure firms (statistically significant at the 5% level with $t = 2.27$), representing an increase of approximately 16% relative to the sample mean of 2.055. Following Diether, Malloy, and Scherbina (2002), who interpret dispersion in analysts' earnings forecasts as a proxy for differences of opinion about a stock's value, and Johnson (2004), who show that dispersion proxies for unpriced information risk arising when asset values are unobservable, this result confirms that the generative AI shock generated genuine disagreement among informed market participants about the future earnings of high-exposure firms — precisely what parameter uncertainty about adoption costs, success probabilities, and competitive consequences predicts.

Panel B tests the left tail prediction. Column (1) shows that deep out-of-the-money put implied volatility (delta $= -20$) rises by 1.9 percentage points for high-exposure firms (statistically significant at the 1% level with $t = 4.03$), confirming that option markets assign meaningfully higher probability to extreme negative outcomes for these firms following ChatGPT's release, consistent with Proposition 3. This result aligns with Drechsler (2013), who models uncertainty as a stochastic process that steepens the implied volatility smirk, and with Miao and Wang (2011), who show that parameter uncertainty fattens the left tail of the asset value distribution by preventing optimal exercise of real options — precisely the mechanism underlying Proposition 3.

In the model, the upside from adoption is bounded by $\delta \cdot w \cdot \lambda / r$ while the downside from failed adoption combined with competitive displacement is unbounded, fattening the left tail disproportionately. Columns (2) and (3) show that the volatility skew and tail risk ratio — both relative measures — are statistically insignificant. This is not inconsistent with Proposition 3: since the vol skew equals OTM put IV minus ATM IV by construction, an insignificant skew coefficient (-0.0012) combined with significant increases in both ATM IV ($+0.0206$) and deep OTM put IV

(+0.0193) implies that the entire implied volatility surface shifts upward for high-exposure firms, with both measures rising by similar magnitudes in absolute terms. The left tail fattens in levels rather than in slope — consistent with a broadening of the risk-neutral distribution rather than a purely asymmetric downside shock.

Table A4 in the Appendix provides complementary evidence for Propositions 2 and 3 when we split the sample along several dimensions. The baseline credit spread effect is concentrated in firms with above-median analyst forecast dispersion (14.9 basis points, $t = 3.18$) and is insignificant for low-dispersion firms (1.7 basis points, $t = 0.89$) — consistent with dispersion proxying for the parameter uncertainty that drives the credit spread effect. Similarly, the effect is nearly twice as large for firms with above-median pre-treatment volatility skew (9.7 basis points, $t = 2.97$) relative to low-skew firms (4.3 basis points, insignificant), confirming that firms with more pre-existing left tail risk face amplified credit spread effects.

5.3 Cross-sectional predictions

Proposition 4 predicts that the spread effect is amplified for firms with greater credit fragility, since these firms sit closer to the default boundary and face disproportionately greater downside risk when parameter uncertainty widens the asset value distribution. Proposition 5 predicts that the effect is moderated for firms with stronger governance, since effective monitoring constrains opportunistic adoption decisions and disciplines the deployment of AI-generated free cash flows. Table 8 tests both predictions.

Columns (1)–(3) test Proposition 4 using triple interaction specifications. Column (1) shows that the spread widening is significantly amplified for firms with lower S&P credit ratings: a one-standard-deviation increase in the rating variable (corresponding to approximately 2.7 notches toward lower credit quality, or roughly one full letter grade) is associated with 9.3 basis points of additional spread widening per standard deviation of GenAI exposure ($t = 4.41$). This result is consistent with Campbell and Taksler (2003), who document that idiosyncratic risk has a larger impact on bond yields for lower-rated firms, and with the convexity of the Merton credit spread in volatility near the default boundary.

Column (2) shows that the effect is amplified by equity return volatility: a one-standard-deviation increase in volatility (0.127) is associated with 9.8 basis points of additional spread widening ($t = 3.90$), consistent with more volatile firms sitting closer to the default boundary and therefore being more sensitive to any increase in asset volatility. Column (3) uses distance to default as a measure of credit fragility: the negative coefficient of −8.6 basis points ($t = -3.64$) confirms that the effect is largest for firms closest to default — a one-standard-deviation decrease in distance to default amplifies the spread widening by 8.6 basis points, consistent with Proposition 4.

Taken together, the triple interaction results in Table 8 imply that the 6.2 basis point baseline effect masks substantial heterogeneity. Table A5 in the Appendix confirms these results when we split our sample by ratings, equity volatility, and distance to default rather than using the triple interaction effect. The effect is concentrated entirely in below-median firms for S&P rating (13.3 basis points, $t = 4.38$, vs −2.4 basis points insignificant), equity volatility (11.1 basis points, $t = 3.53$, vs −1.3 basis points insignificant), and distance to default (15.1 basis points, $t = 4.34$, vs −2.9 basis points insignificant), with the effect statistically indistinguishable from zero in the safe half of the sample in all three cases.

The economic magnitudes for the results in columns (1)–(3) are large. For a firm with credit quality one standard deviation below the sample mean — corresponding to approximately 2.7 notches toward lower ratings, one standard deviation higher equity volatility, or one standard deviation closer to default — the AI-driven spread premium is 12.1, 13.3, and 14.9 basis points respectively, representing 8–10% of the sample mean spread of 148.7 basis points and roughly four times the effect at the average firm. For a firm simultaneously one standard deviation below the mean on all three credit fragility dimensions, the model-implied spread widening reaches approximately 30 basis points — more than ten times the baseline effect and 20% of the sample mean spread. This amplification is consistent with the convexity of the Merton credit spread in volatility near the default boundary: as firms approach financial distress, the same parameter uncertainty shock generates disproportionately larger spread responses, precisely as Campbell and Taksler

(2003) and Collin-Dufresne, Goldstein, and Martin (2001) document for idiosyncratic volatility shocks more broadly.

Columns (4)–(6) test Proposition 5. Column (4) shows that institutional ownership stability moderates the spread widening: a one-standard-deviation increase in IO stability (0.739) reduces the AI-driven spread premium by 7.1 basis points ($t = -3.78$). Institutional investors with stable, long-term holdings have stronger incentives to monitor management and constrain opportunistic AI adoption decisions — for example, restructuring primarily to justify layoffs rather than to maximize long-run productivity — reducing the agency-related component of the credit spread premium. This is consistent with Brockman, Martin, and Unlu (2010) and the broader literature linking institutional ownership to lower cost of debt (Coles, Daniel, and Naveen, 2006). The economic magnitude is substantial: a one-standard-deviation increase in IO stability reduces the AI-driven spread premium by 7.1 basis points — nearly eliminating the baseline effect of 6.2 basis points.

Columns (5) and (6) examine CEO portfolio delta using a split-sample design. The effect is 10.2 basis points for low-delta CEOs ($t = 2.42$) but a statistically insignificant 2.5 basis points for high-delta CEOs ($t = 1.45$). CEO equity alignment — higher sensitivity of pay to firm stock performance — constrains risk-taking behavior and value-extracting adoption decisions, consistent with Coles, Daniel, and Naveen (2006), who show that managerial incentives affect the maturity structure and cost of corporate debt. Table A6 in the Appendix confirms the split-sample pattern: the effect is concentrated in firms with below-median IO stability (10.0 basis points, $t = 3.15$).

Overall, the results in Tables 6, 7, and 8 confirm all five propositions of the model. Equity values rise while credit spreads widen for the same firms exposed to the same shock (Proposition 1). Both implied equity volatility and analyst forecast dispersion rise for high-exposure firms, confirming that option markets immediately price, and financial analysts reflect, the parameter uncertainty surrounding the AI restructuring option (Proposition 2). The increase in risk-neutral volatility is concentrated in the left tail, as reflected in higher deep out-of-the-money put implied volatility (Proposition 3). The spread widening is amplified for firms with greater credit fragility — reaching more than ten times the baseline effect for firms simultaneously distressed on all three dimensions — consistent with the convexity of the Merton credit spread near the default boundary (Proposition 4). Finally, the effect is moderated for firms with strong corporate governance, with the AI-driven spread premium essentially eliminated for firms with stable institutional ownership or high CEO equity alignment (Proposition 5). Parameter uncertainty surrounding a newly created real option is therefore not merely a theoretical construct — it is priced across the capital structure in ways that are quantitatively consistent with the model and cross-sectionally heterogeneous in the directions the model predicts.

5.4 Alternative explanations

We consider and rule out several alternative explanations for the relative credit spread widening. First, the spread increase could reflect general economic uncertainty for firms in technology-adjacent industries rather than parameter uncertainty about a specific real option. Any shock that raises firm-level or industry-level uncertainty could widen credit spreads regardless of the underlying mechanism. The results in Panel C of Table 4, which include industry × time fixed effects that absorb all monthly variation in uncertainty common to firms within the same industry, leave the coefficient on Post × GenAI Exp economically meaningful and statistically significant at 7.2–7.4 basis points.

Second, the widening of the spread could reflect competitive disruption risk rather than a real option held by high-exposure firms. This is the channel in Knesl (2023) in which exposed firms face disruption by more innovative rivals. However, this interpretation predicts falling equity valuations for high-exposure firms, which is directly at odds with the rising equity valuations we document in Table 6: firms that experience a relative increase in credit spreads are the same firms whose equity market valuations rose most sharply following ChatGPT's release, which is inconsistent with a pure disruption story.

Third, the relative widening of spreads could reflect deteriorating bond market liquidity rather than a fundamental repricing of credit risk. The results in columns (6)–(7) of Table 6 rule this out directly, as the bond–CDS basis, which isolates liquidity from default risk, is uniformly small and statistically insignificant at all tenors, while the CDS spread itself widens significantly at the 1- and 3-year horizons, confirming that our results reflect an increase in perceived default risk rather than any illiquidity risk.

A fourth concern is that high-exposure firms may have strategically increased their leverage or accelerated debt issuance post-ChatGPT to fund AI-related capital expenditures, widening spreads through a mechanical increase in leverage rather than any repricing of risk. Evidence in a companion paper directly refutes this. Using the same experimental design, we find that the coefficient on Post$\times$GenAI Exp generates a coefficient of -0.009 ($t = -3.15$) when net debt issuance is the dependent variable, representing approximately 180% of the sample mean — indicating that high-exposure firms substantially reduced debt issuance following ChatGPT's release. The credit spread widening documented above therefore cannot be attributed to an increase in leverage.

Finally, the persistence of the credit spread effect across five successive AI model releases documented in Table 3 rules out a transient investor attention or sentiment shock as the primary driver of our results. For example, if the spread widening reflected a one-time investor reallocation from bonds to equities of high-exposure firms, the effect would moderate as investor attention and reaction faded with each new model release. Taken together, the evidence is consistent with a structural repricing of credit risk driven by parameter uncertainty surrounding a newly created real option, and not with any of the alternative mechanisms above.

6 Conclusion

This paper documents a puzzle in financial markets: firms with high workforce exposure to generative AI experience relatively higher equity valuations alongside relatively higher borrowing costs. Shareholders and bondholders within the same firm, reacting to the same technology shock, reached opposite conclusions regarding firm risk. We argue that this equity-debt divergence is a consequence of these two claimants' asymmetric response to parameter uncertainty surrounding a newly created real option. Specifically, when the parameters of the workforce restructuring real option are uncertain, the distribution of possible asset value outcomes widens. Equity, as a convex claim with downside protection, benefits. Debt, as a concave claim, is sensitive to the now-expanded downside risk and prices it accordingly. A real options model of workforce restructuring with parameter uncertainty accounts for this divergence, quantitatively matching the magnitude and cross-sectional pattern of our empirical estimates.

Our findings have broader implications for the pricing of technology shocks across the capital structure. Unlike the well-specified labor-technology substitution real option in Zhang (2019) with known parameters, which functions as a risk-reducing hedge for all claimants, generative AI introduces parameter uncertainty. The resulting differential between equity and debt investors reflects the asymmetric impact of this uncertainty on their respective claims, rather than any disagreement about AI's potential to restructure the firm's workforce. Parameter uncertainty, therefore, becomes a first-order concern for bondholders, and determines whether a technological shock reduces or raises relative borrowing costs. Our framework provides a tractable way to assess which regime prevails and derives its cross-sectional implications across the capital structure.

These results carry implications for regulators, market participants, and managers. For regulators, treating AI exposure as uniformly positive for all investor classes may understate credit risk for high-exposure firms. For bond investors, the spread differential between high- and low-exposure firms is concentrated in firms with high labor intensity and weak governance — characteristics not typically captured in standard credit risk models. Finally, for managers, the risks associated with AI are priced differently across the capital structure. Bondholders price this risk differently from equity investors, with direct consequences for cost of capital and the firm's ability to reach its optimal capital mix.

Our analysis suggests several directions for future research. While our identification strategy exploits ChatGPT's release as a discrete information event, it captures only a static snapshot of the initial shock and cannot speak to how the real option's value evolves as the technology matures. As adoption costs, success probabilities, and competitive dynamics become better understood, the parameter uncertainty mechanism we document may dissipate, potentially allowing the substitution channel of Zhang (2019) to dominate. Testing this prediction requires a longer time series than is currently available. In addition, our empirical setting is restricted to U.S. capital markets. Whether the documented equity–debt divergence generalizes to other markets — with different institutional frameworks, legal systems, and labor market conditions — remains an open empirical question.

References

Andreadis, L., M. Chatzikonstantinou, E. Kalotychou, C. Louca, and C. Makridis. 2025. Local Heterogeneity in Artificial Intelligence Jobs Over Time and Space. *AEA Papers and Proceedings* 115:29–34.

Andrews, I., and M. Farboodi. 2025. Do Markets Believe in Transformative AI? Working Paper 34243, National Bureau of Economic Research.

Babina, T., A. Fedyk, A. X. He, and J. Hodson. 2024. Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics* 151:103745.

Babina, T., A. Fedyk, A. X. He, and J. Hodson. 2025. Artificial Intelligence and Firms' Systematic Risk. *SSRN Working Paper* .

Baker, A. C., D. F. Larcker, and C. C. Y. Wang. 2022. How Much Should We Trust Stacked-Difference-in-Differences Estimates? *Review of Financial Studies* 35:3344–3390.

Baker, S. R., N. Bloom, and S. J. Davis. 2016. Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics* 131:1593–1636.

Becker, B., and V. Ivashina. 2023. Disruption and Credit Markets. *Journal of Finance* 78:105–139.

Brockman, P., X. Martin, and E. Unlu. 2010. Executive Compensation and the Maturity Structure of Corporate Debt. *Journal of Finance* 65:1123–1161.

Campbell, J. Y., and G. B. Taksler. 2003. Equity Volatility and Corporate Bond Yields. *Journal of Finance* 58:2321–2350.

Chen, L., D. A. Lesmond, and J. Wei. 2007. Corporate Yield Spreads and Bond Liquidity. *Journal of Finance* 62:119–149.

Chen, Q., and Y. Han. 2023. Options Market Ambiguity and Its Information Content. *Journal of Financial Markets* 64:100790.

CNBC. 2025. Why AI May Kill Career Advancement for Many Young Workers. URL https://www.cnbc.com/2025/11/20/why-ai-may-kill-career-advancement-for-many-young-workers.html.

Coles, J. L., N. D. Daniel, and L. Naveen. 2006. Managerial Incentives and Risk-Taking. *Journal of Financial Economics* 79:431–468.

Collin-Dufresne, P., R. S. Goldstein, and J. S. Martin. 2001. The Determinants of Credit Spread Changes. *Journal of Finance* 56:2177–2207.

Crunchbase. 2025. Tech Layoffs: US Companies with Job Cuts in 2024 and 2025. URL https://news.crunchbase.com/startups/tech-layoffs/.

Dick-Nielsen, J. 2009. Liquidity Biases in TRACE. *Journal of Fixed Income* 19:43–55.

Dick-Nielsen, J. 2014. How to Clean Enhanced TRACE Data. *SSRN Working Paper* .

Dickerson, A., C. Robotti, and G. Rossetti. 2023. Common Pitfalls in the Evaluation of Corporate Bond Strategies. Tech. rep., UNSW Business School.

Diether, K. B., C. J. Malloy, and A. Scherbina. 2002. Differences of Opinion and the Cross Section of Stock Returns. *Journal of Finance* 57:2113–2141.

Dixit, A. K., and R. S. Pindyck. 1994. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.

Donangelo, A. 2014. Labor Mobility: Implications for Asset Pricing. *Journal of Finance* 69:1321–1346.

Drechsler, I. 2013. Uncertainty, Time-Varying Fear, and Asset Prices. *Journal of Finance* 68:1843–1889.

Eisfeldt, A. L., G. Schubert, and M. B. Zhang. 2026. Generative AI and Firm Values. *Journal of Finance*, forthcoming.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock. 2024. GPTs are GPTs: Labor Market Impact Potential of Large Language Models. *Science* 384:1306–1308.

Elyasiani, E., and J. Jia. 2010. Distribution of Institutional Ownership and Corporate Firm Performance. *Journal of Banking & Finance* 34:606–620.

Epstein, L. G., and M. Schneider. 2010. Ambiguity and Asset Markets. *Annual Review of Financial Economics* 2:315–346.

Fama, E. F., and K. R. French. 1997. Industry Costs of Equity. *Journal of Financial Economics* 43:153–193.

Faulkender, M., and M. A. Petersen. 2006. Does the Source of Capital Affect Capital Structure? *Review of Financial Studies* 19:45–79.

Gilchrist, S., and E. Zakrajšek. 2012. Credit Spreads and Business Cycle Fluctuations. *American Economic Review* 102:1692–1720.

Goodman-Bacon, A. 2021. Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics* 225:254–277.

Gürkaynak, R. S., B. Sack, and J. H. Wright. 2007. The U.S. Treasury Yield Curve: 1961 to the Present. *Journal of Monetary Economics* 54:2291–2304.

Hampole, M., D. Papanikolaou, L. D. W. Schmidt, and B. Seegmiller. 2025. Artificial Intelligence and the Labor Market. Working Paper 33509, National Bureau of Economic Research.

Jensen, M. C. 1986. Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers. *American Economic Review* 76:323–329.

Johnson, T. C. 2004. Forecast Dispersion and the Cross-Section of Expected Returns. *Journal of Finance* 59:1957–1986.

Kaviani, M. S., L. Kryzanowski, H. Maleki, and P. Savor. 2020. Policy Uncertainty and Corporate Credit Spreads. *Journal of Financial Economics* 138:838–865.

Knesl, J. 2023. Automation and the Displacement of Labor by Capital: Asset Pricing Theory and Empirical Evidence. *Journal of Financial Economics* 147:271–296.

Kogan, L., and D. Papanikolaou. 2014. Growth Opportunities, Technology Shocks, and Asset Prices. *Journal of Finance* 69:675–718.

Kogan, L., D. Papanikolaou, L. D. W. Schmidt, and B. Seegmiller. 2023. Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data. Working Paper 31846, National Bureau of Economic Research.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman. 2017. Technological Innovation, Resource Allocation, and Growth. *Quarterly Journal of Economics* 132:665–712.

Kogan, L., D. Papanikolaou, and N. Stoffman. 2020. Left Behind: Creative Destruction, Inequality, and the Stock Market. *Journal of Political Economy* 128:855–906.

Layoffs.fyi. 2025. Tech and Startup Layoff Tracker. URL https://layoffs.fyi.

Liu, Y., and J. C. Wu. 2021. Reconstructing the Yield Curve. *Journal of Financial Economics* 142:1395–1425.

Longstaff, F. A., S. Mithal, and E. Neis. 2005. Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit Default Swap Market. *Journal of Finance* 60:2213–2253.

McDonald, R., and D. Siegel. 1986. The Value of Waiting to Invest. *Quarterly Journal of Economics* 101:707–727.

Merton, R. C. 1974. On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance* 29:449–470.

Miao, J., and N. Wang. 2011. Risk, Uncertainty, and Option Exercise. *Journal of Economic Dynamics and Control* 35:442–461.

Novy-Marx, R. 2011. Operating Leverage. *Review of Finance* 15:103–134.

Papanikolaou, D. 2011. Investment Shocks and Asset Prices. *Journal of Political Economy* 119:639–685.

Pástor, v., and P. Veronesi. 2009. Technological Revolutions and Stock Prices. *American Economic Review* 99:1451–1483.

Serfling, M. 2016. Firing Costs and Capital Structure Decisions. *Journal of Finance* 71:2239–2286.

Simon, L. 2025. Entry-Level Job Postings and AI Exposure. URL https://www.reveliolabs.com/news/macro/is-ai-responsible-for-the-rise-in-entry-level-unemployment/.

Veronesi, P. 1999. Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model. *Review of Financial Studies* 12:975–1007.

Zhang, M. B. 2019. Labor-Technology Substitution: Implications for Asset Pricing. *Journal of Finance* 74:1793–1839.

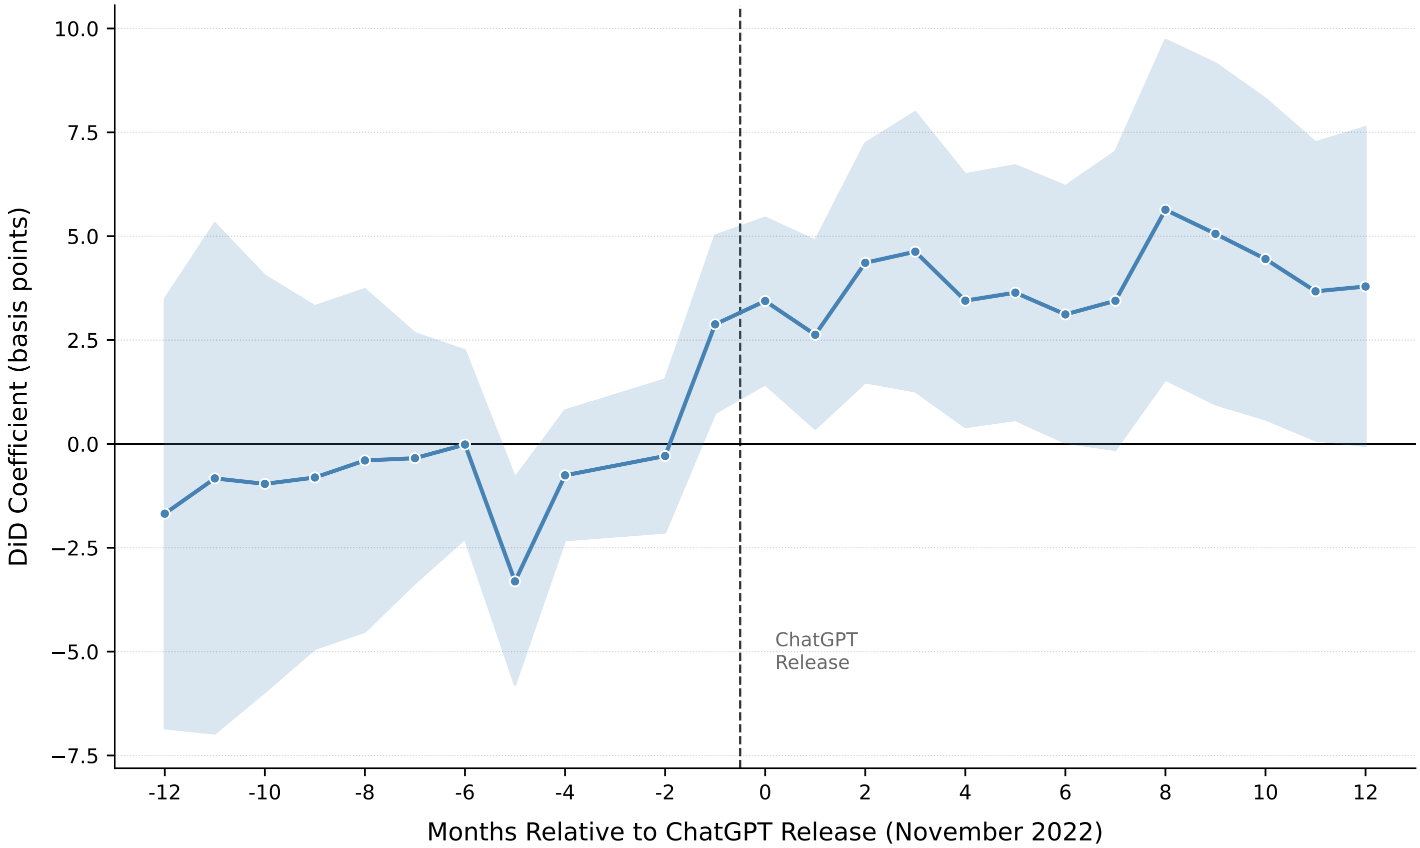

Figure 1: Parallel trends: Generative AI exposure and credit spreads around ChatGPT's release

Notes: This figure plots the monthly event-study coefficients $\hat{\beta}_\tau$ from the regression $\text{Spread}_{i,b,t} = \sum_{\tau \neq -3} \beta_\tau \cdot \mathbf{1}[t = \tau] \times \text{GenAI Exp}_i + \alpha_b + \gamma_t + \mathbf{X}'_{i,b,t}\delta + \varepsilon_{i,b,t}$, where τ denotes months relative to ChatGPT's release in November 2022 (month 0). The dependent variable is the credit spread in basis points. GenAI exposure is standardized to mean zero and unit standard deviation. The regression includes bond fixed effects, year-month fixed effects, and the full set of controls from Table 2. The shaded region is the 95% confidence interval based on standard errors clustered at the firm level. Month -3 is omitted. The vertical dashed line marks November 2022.

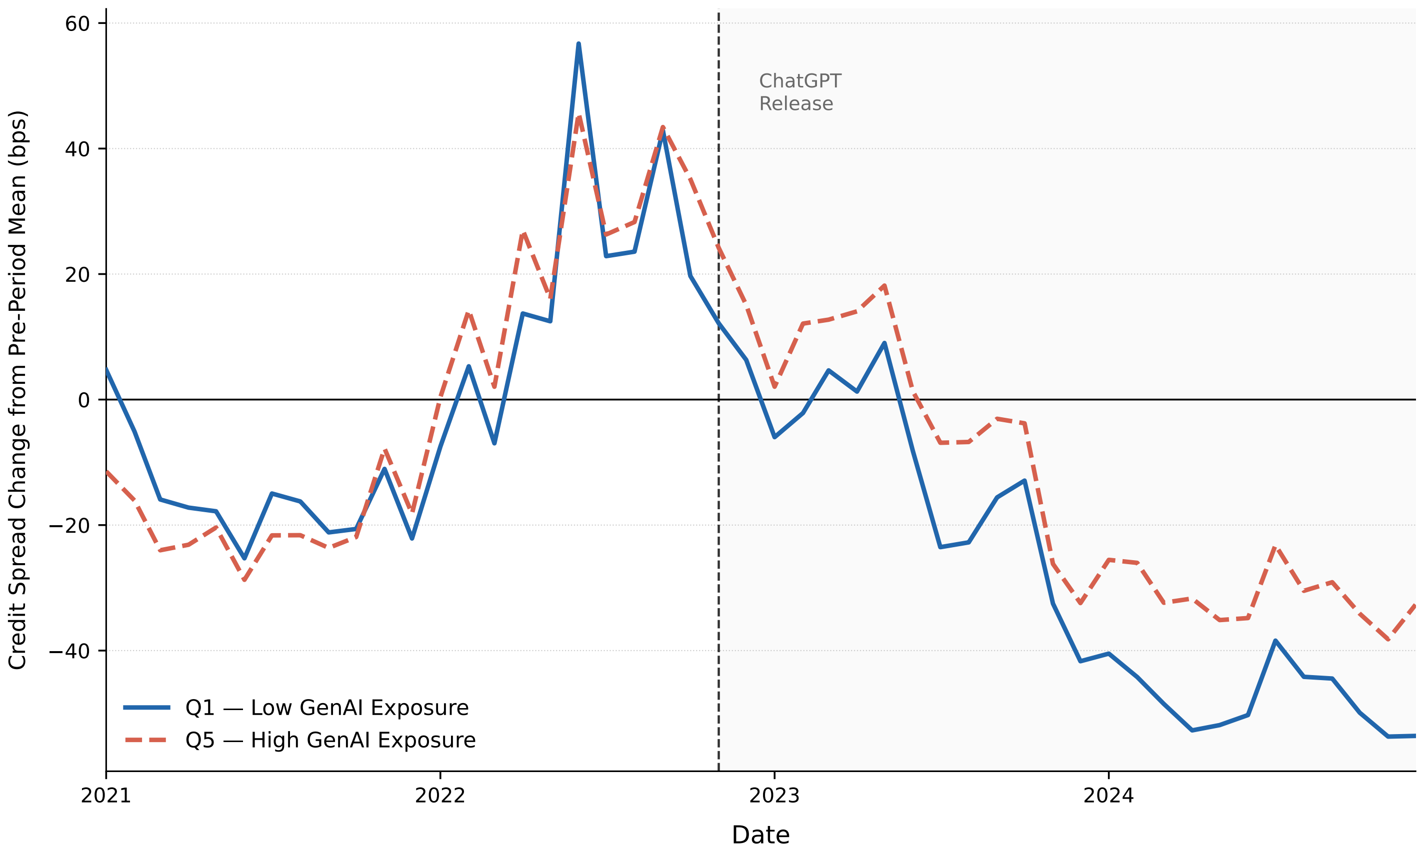

Figure 2: Credit spread dynamics by generative AI exposure quintile

Notes: This figure plots the monthly value-weighted average credit spread for the lowest-exposure quintile (Q1, solid blue) and the highest-exposure quintile (Q5, dashed red), demeaned by each quintile's pre-ChatGPT mean computed over January 2021 through October 2022. Firms are sorted into quintiles annually based on their GenAI exposure score from Eisfeldt, Schubert, and Zhang (2026). The credit spread is the bond yield-to-maturity minus the matched-maturity Treasury yield interpolated from the Liu and Wu (2021) yield curve, averaged across bonds within each firm using face-value weights and then averaged across firms within each quintile using market capitalization weights. The vertical dashed line marks November 2022, the month of ChatGPT's public release.

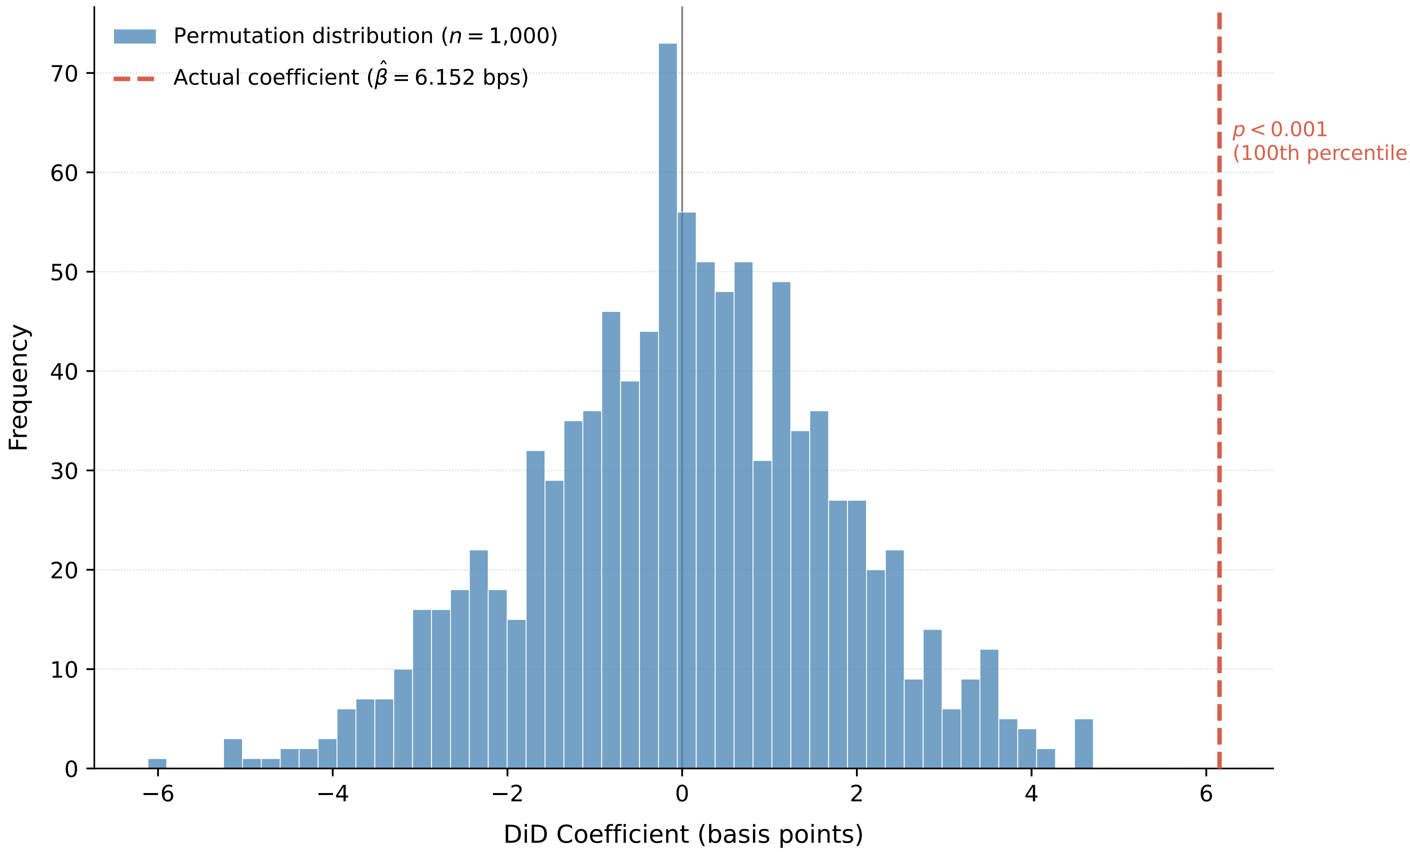

Figure 3: Fisher randomization inference

Notes: This figure plots the distribution of 1,000 difference-in-differences coefficients obtained by randomly reassigning GenAI exposure scores across firms while holding the panel structure, fixed effects, and controls fixed. The dependent variable is the monthly volume-weighted credit spread in basis points. Each permutation uses the preferred specification from Table 2 (bond and year-month fixed effects). The vertical dashed line marks the actual estimated coefficient of 6.15 basis points from column (2) of Table 2.

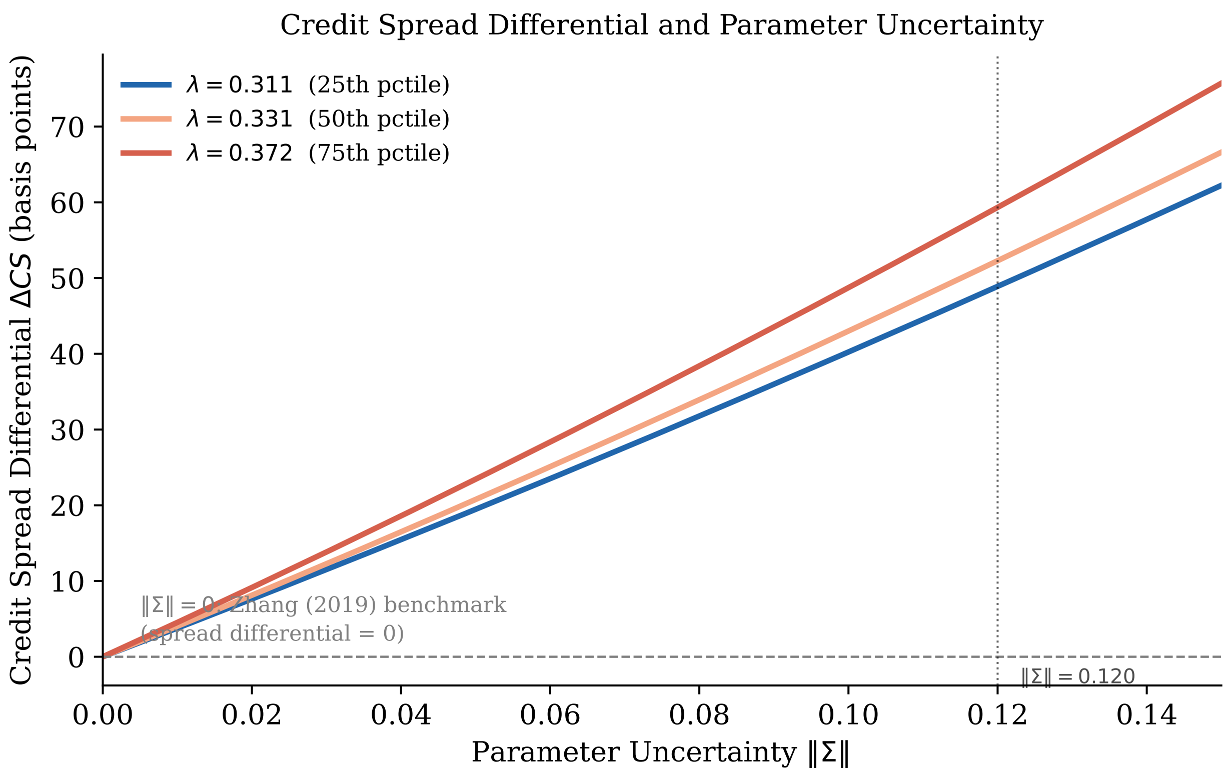

Figure 4: Comparative statics: Credit spread differential and parameter uncertainty

Notes: This figure plots the credit spread differential ΔCS (in basis points) as a function of parameter uncertainty $\|\Sigma\|$ for three levels of generative AI workforce exposure λ. The credit spread differential for a firm with workforce exposure ($\lambda > 0$) is benchmarked to the spread for a baseline firm with no workforce exposure ($\lambda = 0$), both computed using the model with calibrated parameters as reported in Table 5. Exposure levels correspond to the 25^{th}, 50^{th}, and 75^{th} percentiles of the Eisfeldt, Schubert, and Zhang (2026) workforce exposure distribution.

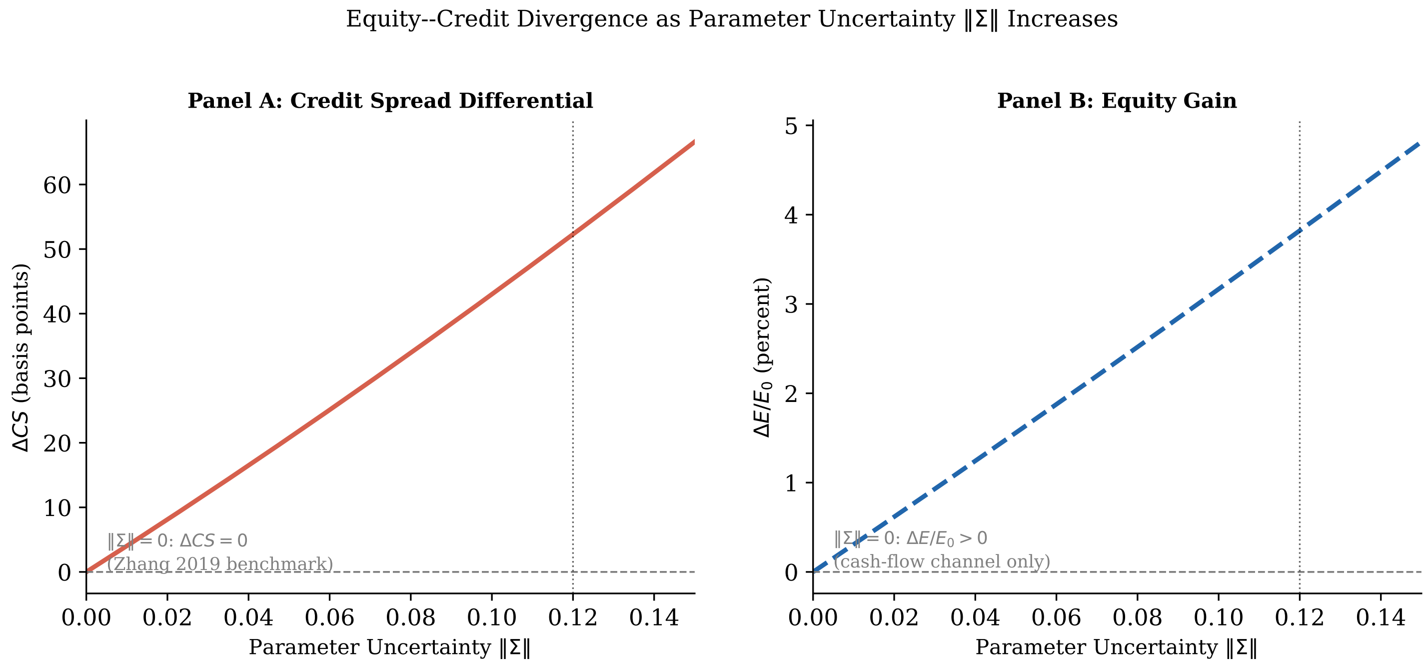

Figure 5: Comparative statics: Equity–Credit divergence and parameter uncertainty

Notes: This figure plots the credit spread differential ΔCS in basis points and equity gain $\Delta E/E_0$ as a function of parameter uncertainty $\|\Sigma\|$ for a firm with median generative AI workforce exposure ($\lambda = 0.331$, corresponding to the 50^{th} percentile of the Eisfeldt, Schubert, and Zhang (2026) distribution). Panel A plots ΔCS, defined as the credit spread for a firm with workforce exposure ($\lambda > 0$) benchmarked to the spread for a baseline firm with no workforce exposure ($\lambda = 0$) computed using the model with calibrated parameters as reported in Table 5. Panel B plots $\Delta E/E_0$ for the same set of firms.

Table 1: Summary statistics

Notes: This table reports summary statistics for the full sample. Panel A reports statistics for the dependent variable, Spread (basis points), computed as the monthly volume-weighted average bond yield minus the zero-coupon Treasury yield at the bond's remaining maturity, interpolated from the Liu and Wu (2021) yield curve. Panel B reports statistics for GenAI exposure, the firm-level measure of workforce exposure to generative AI constructed by Eisfeldt, Schubert, and Zhang (2026) based on occupation-level employment shares from Revelio Labs and the task-exposure rubric of Eloundou, Manning, Mishkin, and Rock (2024). GenAI exposure is standardized with mean zero and unit standard deviation. Panel C reports statistics for firm-level variables. Panel D reports statistics for bond-level variables. Variable definitions are in Table A1 in Appendix A1. Monthly data 2021 – 2024.

	Mean	σ	25^{th}	Median	75^{th}
Panel A: Credit spreads					
Spread	148.663	113.019	78.701	117.986	179.407
Spread over AAA	86.543	116.590	22.036	71.452	126.399
Log spread	4.771	0.684	4.366	4.771	5.190
Panel B: Generative AI exposure					
GenAI exp (raw)	0.341	0.055	0.311	0.331	0.372
GenAI exp (std)	-0.240	0.932	-0.742	-0.407	0.291
Panel C: Firm level variables					
Assets	10.883	1.273	10.031	10.978	11.812
Leverage	0.404	0.162	0.292	0.409	0.476
Tangibility	0.396	0.274	0.135	0.348	0.664
Liquidity	0.071	0.077	0.012	0.046	0.112
Tobin's Q	1.985	1.550	1.186	1.392	2.236
ROA	0.055	0.065	0.023	0.041	0.083
Return	0.011	0.091	-0.043	0.008	0.036
Volatility	0.293	0.127	0.208	0.258	0.345
Cash flow vol	0.028	0.028	0.010	0.018	0.036
Dispersion	2.055	3.122	1.092	1.508	1.870
CEOΔ	6.126	1.561	5.178	6.168	7.164
Institutional	2.475	0.739	1.972	2.423	2.911
Panel D: Bond-level variables					
Coupon	4.452	1.627	3.375	4.300	5.450
Maturity	11.369	9.940	4.085	7.384	18.253
Volume	13.074	1.160	12.766	13.122	13.710
Rating	8.426	2.707	7.000	8.000	10.000

Table 2: Generative AI exposure and credit spreads: Baseline difference-in-differences

Notes: This table reports difference-in-differences estimates of the effect of generative AI exposure on corporate bond credit spreads. The dependent variable is the monthly volume-weighted average credit spread in basis points, measured as the bond yield-to-maturity minus the matched-maturity Treasury yield interpolated from the Liu and Wu (2021) yield curve. The key independent variable, Post $\times$ GenAI Exp, interacts an indicator equal to one for year $\geq$ 2023 with the firm-level workforce exposure to generative AI from Eisfeldt, Schubert, and Zhang (2026), standardized to mean zero and unit standard deviation. Controls include assets (log), ROA, leverage, Tobin's Q, tangibility, liquidity, return, cash flow volatility (5-year rolling), remaining maturity, and trading volume (log). The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)	(3)	(4)
Post $\times$ GenAI Exp	7.928***	6.152***	6.232***	7.320***
	(3.63)	(2.99)	(3.03)	(3.46)
Assets	-20.489	-2.247	-2.306	-10.344
	(-1.56)	(-0.22)	(-0.23)	(-1.08)
ROA	-37.656	-90.181***	-88.388***	-84.315***
	(-1.09)	(-3.08)	(-3.02)	(-2.74)
Leverage	189.260***	194.541***	195.113***	207.688***
	(5.86)	(6.13)	(6.15)	(6.51)
Tobin's Q	-11.294***	-2.679	-2.526	-3.341
	(-2.93)	(-1.17)	(-1.11)	(-1.45)
Tangibility	26.233	73.050	72.789	50.756
	(0.41)	(1.33)	(1.34)	(0.97)
Liquidity	27.631	113.679***	114.681***	116.133***
	(0.67)	(3.08)	(3.12)	(3.03)
Return	-39.053***	-13.181***	-26.855***	-19.773***
	(-10.57)	(-3.33)	(-7.70)	(-4.60)
Cash flow vol	189.067***	103.668	100.315	102.277
	(3.91)	(1.55)	(1.45)	(1.15)
Maturity	11.596***	-59.850***	12.202***	2.425***
	(10.73)	(-4.90)	(12.82)	(25.25)
Volume (log)	12.590*	20.456***	20.141***	-6.319***
	(1.93)	(2.79)	(2.83)	(-6.13)
Controls	Yes	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes	No
Firm FE	No	No	No	Yes
Year-Month FE	No	Yes	No	No
Year FE	No	No	Yes	Yes
Adj. R^2	0.839	0.878	0.868	0.750
Observations	308,248	308,248	308,248	308,248

Table 3: Multi-event analysis

Notes: This table reports multi-event difference-in-differences estimates validating the baseline result across major generative AI releases. The dependent variable is the monthly volume-weighted credit spread in basis points throughout. Panel A reports stacked difference-in-differences estimates following Baker, Larcker, and Wang (2022) across five major GenAI events: ChatGPT (November 2022), GPT-4 (March 2023), Llama 2 (July 2023), Gemini (December 2023), and GPT-4o (May 2024). Each event defines a cohort; stacking across cohorts avoids contamination from already-treated observations serving as controls in a standard two-way fixed effects design. Panel B replaces the binary Post indicator with the count of cumulative GenAI events that have occurred up to each month, interacted with firm-level GenAI exposure. The cumulative count equals zero before November 2022, one after ChatGPT's release, two after GPT-4, and so on up to five after GPT-4o. The coefficient measures the average incremental effect on credit spreads of each additional GenAI event per standard deviation of firm exposure. Controls and fixed effects follow Table 2 except where noted. The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)	(3)	(4)
Panel A: Stacked DiD across five GenAI events				
Post × GenAI Exposure	1.425***	1.070**	1.531***	2.366***
	(2.87)	(2.08)	(3.10)	(4.69)
Controls	Yes	Yes	Yes	Yes
Cohort × Bond FE	Yes	Yes	Yes	No
Cohort × Firm FE	No	No	No	Yes
Cohort × Year-Month FE	No	Yes	No	No
Cohort × Year FE	No	No	Yes	Yes
Adj. R^2	0.947	0.953	0.949	0.800
Observations	222,871	222,871	222,871	224,392
Panel B: Cumulative GenAI events × exposure				
Cumul. Events × GenAI Exp	2.708***	1.564***	1.742***	2.009***
	(3.80)	(2.75)	(3.00)	(3.46)
Controls	Yes	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes	No
Firm FE	No	No	No	Yes
Year-Month FE	No	Yes	No	No
Year FE	No	No	Yes	Yes
Adj. R^2	0.840	0.878	0.868	0.750
Observations	308,248	308,248	308,248	308,248

Table 4: Robustness

Notes: This table reports robustness checks for the baseline difference-in-differences result. The key independent variable is Post × GenAI Exp throughout. Panel A excludes transition months around ChatGPT's release, dropping November–December 2022 in column (1) and November 2022 through January 2023 in column (2), using bond and year-month fixed effects. Panel B uses alternative credit spread measures as the dependent variable: the Gilchrist and Zakrajšek (2012) (GZ spread) in column (1) and the spread over AAA-rated non-financial corporate bond yields in column (2), using bond and year-month fixed effects. Panel C aggregates bonds to the firm-month level using face-value weights and augments the specification with industry-specific time fixed effects: firm and year-month fixed effects in column (1), firm and Fama and French (1997) 49-industry × year fixed effects in column (2), and firm and FF49-industry × year-month fixed effects in column (3). Panel D restricts the sample to non-callable bonds, using the same four fixed effect structures as Table 2. Controls follow Table 2 throughout. The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)	(3)	(4)
Panel A: Excluding transition months				
Post × GenAI Exposure	6.588***	6.573***		
	(3.05)	(3.00)		
Controls	Yes	Yes		
Bond FE	Yes	Yes		
Year-Month FE	Yes	Yes		
Sample	Drop Nov–Dec 2022	Drop Nov22–Jan23		
Adj. R^2	0.878	0.878		
Observations	295,597	289,320		
Panel B: Alternative dependent variables				
Post × GenAI Exposure	5.023***	5.678**		
	(2.73)	(2.19)		
Controls	Yes	Yes		
Bond FE	Yes	Yes		
Year-Month FE	Yes	Yes		
Dep. variable	GZ Spread	AAA Spread		
Adj. R^2	0.891	0.856		
Observations	308,248	308,248		
Panel C: Firm-month level + Industry × Time fixed effects				
Post × GenAI Exposure	10.932***	7.399**	7.248**	
	(4.41)	(2.34)	(2.24)	
Controls	Yes	Yes	Yes	
Firm FE	Yes	Yes	Yes	
Year-Month FE	Yes	No	No	
FF49 × Year FE	No	Yes	No	
FF49 × Year-Month FE	No	No	Yes	
Adj. R^2	0.872	0.874	0.885	
Observations	33,941	33,941	33,807	
Panel D: Non-callable bond subsample				
Post × GenAI Exposure	22.680***	20.295***	20.176***	21.007***
	(3.77)	(3.90)	(3.87)	(3.85)
Controls	Yes	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes	No
Firm FE	No	No	No	Yes
Year-Month FE	No	Yes	No	No
Year FE	No	No	Yes	Yes
Adj. R^2	0.783	0.830	0.819	0.714
Observations	19,656	19,656	19,656	19,656

Table 5: Calibration parameters

Notes: This table reports the parameter values used to calibrate the theoretical model in Section 4. Baseline asset volatility is set equal to the sample mean equity volatility from Panel C of Table 1, consistent with the Merton (1974) framework in which equity volatility proxies for asset volatility (Campbell and Taksler, 2003). Leverage is the sample mean from Panel C of Table 1 (Collin-Dufresne, Goldstein, and Martin, 2001). Bond maturity is the sample median from Panel D of Table 1. The interquartile range of GenAI exposure is drawn from Panel B of Table 1 and represents the spread between firms at the 75^{th} and 25^{th} percentiles of raw exposure. Parameter uncertainty $\|\Sigma\|$ is the sole free parameter, calibrated to match the baseline DiD estimate of 8 basis points per standard deviation of GenAI exposure reported in Table 2.

	Symbol	Value	Source
Panel A: Market and firm parameters			
Baseline asset volatility	σ_0	0.293	Sample mean equity volatility
Risk-free rate	r	0.040	Sample period average
Leverage ratio	L	0.404	Sample mean
Bond maturity	$\mathcal{T}$	7.4 years	Sample median
Panel B: GenAI exposure			
IQR of GenAI exposure	$\Delta\lambda$	0.061	Eisfeldt, Schubert, and Zhang (2026); Panel B
Panel C: Calibrated parameter			
Parameter uncertainty	$\|\Sigma\|$	0.120	Calibrated to baseline DiD result
Panel D: Model-implied predictions			
Spread differential per SD	ΔCS	8.0 bps	Target (baseline DiD estimate)
Spread differential IQR	ΔCS	3.6 bps	Model-implied
Equity impact per SD	$\Delta E/E_0$	0.35%	Model-implied (volatility effect)
Eisfeldt et al. equity CAR		∼1.7%	Eisfeldt, Schubert, and Zhang (2026)

Table 6: Equity-bond divergence: Proposition 1

Notes: This table tests Proposition 1 by examining the joint response of equity and credit markets to GenAI exposure following ChatGPT's release. The key independent variable is Post × GenAI Exp throughout. Column (1) uses monthly stock return as the dependent variable at the firm-month level with firm and year-month fixed effects. Column (2) uses Tobin's Q at the firm-year level with firm and year fixed effects. Columns (3)–(7) use firm-month data restricted to the CDS subsample with firm and year-month fixed effects. Column (3) uses the value-weighted bond credit spread; columns (4)–(5) use CDS spreads at the 1- and 3-year tenors; columns (6)–(7) use the bond–CDS basis (bond spread minus CDS spread) at the 1- and 3-year tenors, which captures bond-specific liquidity and funding costs beyond pure default risk. CDS data are from IHS Markit, USD-denominated, senior unsecured, composite quotes. Controls follow Table 2 throughout. The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1) Stock Ret.	(2) Tobin's Q	(3) Bond Spread	(4) CDS 1Y	(5) CDS 3Y	(6) Liq. 1Y	(7) Liq. 3Y
Post × GenAI Exp	0.0055***	0.1086**	10.666***	7.372**	8.910**	3.244	1.738
	(4.21)	(2.03)	(3.32)	(2.33)	(2.11)	(1.04)	(0.58)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	No	Yes	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	No	No	No	No
Adj. R^2	0.250	0.857	0.893	0.607	0.693	0.623	0.567
Observations	33,941	2,873	11,374	11,253	11,317	11,253	11,317

Table 7: Volatility channel and left tail risk: Propositions 2 and 3

Notes: This table tests Propositions 2 and 3 using option-implied measures from OptionMetrics. All regressions are at the firm-month level with firm and year-month fixed effects and the full set of controls from Table 2, except column (2) which uses analyst forecast dispersion as the dependent variable at the firm-year level with firm and year fixed effects. Columns (1)–(2) test Proposition 2: column (1) uses ATM implied volatility and column (2) uses analyst forecast dispersion (standard deviation of IBES EPS forecasts scaled by the absolute mean forecast) as dependent variables. Columns (3)–(5) test Proposition 3: column (3) uses deep out-of-the-money put implied volatility (delta = −20), column (4) uses the volatility skew (OTM put IV minus ATM IV), and column (5) uses the tail risk ratio (OTM put IV divided by ATM IV). The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)	(3)	(4)	(5)
	ATM IV	Dispersion	Deep OTM Put IV	Vol Skew	Tail Risk Ratio
	Panel A: Proposition 2		Panel B: Proposition 3		
Post × GenAI Exp	0.0206***	0.3305**	0.0193***	-0.0012	-0.0064
	(5.12)	(2.27)	(4.03)	(-0.45)	(-0.97)
Controls	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	No	Yes	Yes	Yes
Year FE	No	Yes	No	No	No
Adj. R^2	0.714	0.436	0.751	0.348	0.410
Observations	33,427	2,749	33,427	33,427	33,427

Table 8: Cross-sectional predictions: Propositions 4 and 5

Notes: This table tests Propositions 4 and 5. The dependent variable is the monthly volume-weighted credit spread in basis points throughout. Columns (1)–(3) test Proposition 4 using triple interaction specifications with the conditioning variable standardized to mean zero and unit standard deviation: column (1) uses S&P credit rating (higher values indicate lower ratings, i.e., greater credit risk); column (2) uses equity return volatility; column (3) uses distance to default (higher values indicate greater distance from default, i.e., lower credit risk). Columns (4)–(6) test Proposition 5: column (4) uses a triple interaction with institutional ownership stability (Elyasiani and Jia, 2010); columns (5) and (6) split the sample at the median pre-treatment CEO portfolio delta (Coles, Daniel, and Naveen, 2006), with column (5) covering low-delta and column (6) high-delta firms. All regressions include bond and year-month fixed effects and the full set of controls from Table 2. The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)	(3)	(4)	(5)	(6)
	Rating	Volatility	Dist. to Default	Instititional	Low CEOΔ	High CEOΔ
	Panel A: Proposition 4			Panel B: Proposition 5		
Post × GenAI Exp × Cond. Var	9.259***	9.794***	-8.608***	-7.074***	—	—
	(4.41)	(3.90)	(-3.64)	(-3.78)		
Post × GenAI Exp	2.881*	3.505**	6.316***	5.864***	10.187**	2.535
	(1.66)	(2.24)	(3.08)	(3.09)	(2.42)	(1.45)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	Full	Full	Full	Full	Low Delta	High Delta
Adj. R^2	0.886	0.883	0.880	0.881	0.865	0.881
Observations	307,825	307,735	291,986	302,207	133,175	132,217

Appendices

Equity Prices the Opportunity, Debt Prices the Risk: Generative AI and Corporate Credit Spreads.

A1 Variable definitions and additional statistics

Table A1: Variable definitions

Variable	Definition	Source
Spread	Monthly volume-weighted bond yield minus the zero-coupon Treasury yield at the bond's remaining maturity, interpolated from the Liu and Wu (2021) yield curve	TRACE
Spread over AAA	Difference between bond yield to maturity and the yield to maturity on an index of value-weighted AAA-rated bonds issued by U.S. nonfinancial corporations	TRACE
Log spread	Natural logarithm of credit spread	
GenAI exp	Share of a firm's workforce tasks that generative AI can perform or augment, constructed by linking occupation-level employment shares from Revelio Labs to the task-exposure rubric of Eloundou, Manning, Mishkin, and Rock (2024).	Eisfeldt, Schubert, and Zhang (2026)
Post	Binary variable equal to one for observations in 2023–2024 and zero otherwise	
Assets	Natural logarithm of total assets (at)	Compustat

Continued on next page

Variable	Definition	Source
Leverage	Long- plus short-term debt (dltt and dlc) divided by total assets	Compustat
Tangibility	Net property, plant, and equipment (ppent) divided by total assets	Compustat
Liquidity	Cash and short-term investments (che) divided by total assets	Compustat
Tobin's Q	Market value of equity plus book value of total debt, divided by total assets	Compustat
ROA	Income before extraordinary items (ib) divided by total assets	Compustat
Return	Monthly return including dividends (ret)	CRSP
Volatility	Standard deviation of daily excess stock returns over risk-free rate (annualized)	CRSP
Cash flow vol	Standard deviation of firm cash flows, computed over backward-looking 5-year rolling window	Compustat
Dispersion	Standard deviation of individual analyst EPS forecasts divided by the absolute value of the mean forecast, computed from all forecasts issued within 365 days before fiscal year end Diether, Malloy, and Scherbina (2002)	IBES
CEOΔ	Natural logarithm of one plus the CEO's portfolio delta, where delta is the dollar change in the CEO's equity and option holdings for a 1% increase in stock price Coles, Daniel, and Naveen (2006)	ExecuComp

Continued on next page

Variable	Definition	Source
Institutional	For each institutional investor j, $IOP_j = \bar{h}_j/\sigma(h_j)$, where $\bar{h}_j$ and $\sigma(h_j)$ are the mean and standard deviation of investor j's percentage holding over a 20-quarter rolling window. Firm-level institutional ownership stability is the average IOP_j across all institutional holders Elyasiani and Jia (2010)	13F
Coupon	Annual coupon rates in percentages	FISD
Maturity	Number of years remaining until the bond's maturity date	FISD
Volume	Natural logarithm of monthly dollar trading volume	TRACE
Rating	S&P credit rating mapped to a numeric scale: 1 (AAA) to 22 (D)	FISD

Table A2: Sample distribution by industry

Notes: This table reports the distribution of bond-month observations across Fama and French (1997) twelve-industry classifications. Financial firms (SIC codes 6000–6999) are excluded from the sample. Monthly data 2021–2024.

Industry	Bond-Months	Pct (%)
Utilities	66,360	21.5
Other	42,514	13.8
Manufacturing	37,035	12.0
Business Equipment	28,048	9.1
Telephone and TV	24,770	8.0
Healthcare, Drugs	23,211	7.5
Oil, Gas, and Coal	22,849	7.4
Wholesale, Retail	22,462	7.3
Consumer NonDurables	17,711	5.7
Consumer Durables	12,478	4.0
Chemicals	10,810	3.5
Total	308,248	100.0

A2 Model proofs

This appendix provides formal derivations for the five propositions in Section 4. Throughout, we use effective asset volatility $\sigma = \sigma_0 + \lambda\|\Sigma\|$ under $\mathbb{Q}$, the adoption NPV in equation (6), and d_1, d_2 as defined in equation (4). All comparative statics are well-defined for $V_0 > 0$, $F > 0$, $\mathcal{T} > 0$, $r \geq 0$, $\sigma > 0$.

Proof of Proposition 1 (Equity-debt divergence)

Part (i): Cash flow effect ($\mathbb{P}$). Under $\mathbb{P}$, $\mathbb{E}^P[V_\mathcal{T}] = V_0 \cdot \exp([\mu_0 + \Delta\mu]\mathcal{T}) > V_0 \cdot \exp(\mu_0\mathcal{T})$ since $\Delta\mu > 0$. Equity is increasing in expected asset value. □

Part (iii): Volatility effect on credit spreads ($\mathbb{Q}$). Differentiating equation (5) with respect to σ under $\mathbb{Q}$:

$$\frac{\partial(CS)}{\partial\sigma} = \frac{n(d_2) \cdot \sqrt{\mathcal{T}} \cdot Fe^{-r\mathcal{T}}}{\mathcal{T} \cdot D_0} > 0$$

since $n(d_2) > 0$, $\mathcal{T} > 0$, $D_0 > 0$, and $Fe^{-r\mathcal{T}} > 0$. Under $\mathbb{Q}$ the drift μ is replaced by r, so $\Delta\mu$ does not enter equation (5). Credit spreads widen exclusively through the volatility effect. □

Part (iv): Value redistribution. Total firm value $V_0 = E_0 + D_0$ increases because $\Phi > 0$. The equity gain from both effects exceeds the bond loss from the volatility effect alone ($\Delta E > |\Delta D|$), so total value rises while redistribution favors equity. □

Proof of Proposition 2 (Implied equity volatility)

At $t = 0$ the firm has not yet adopted, so current cash flows are unchanged from the pre-AI regime. Under $\mathbb{Q}$, effective volatility immediately becomes $\sigma_0 + \lambda\|\Sigma\|$ because option prices incorporate the full distribution of future adoption outcomes including continuous belief revision (Veronesi, 1999). The cross-sectional prediction follows from $\partial\sigma/\partial\lambda = \|\Sigma\| > 0$: firms with greater AI workforce exposure λ face a larger volatility increase under $\mathbb{Q}$ and therefore higher implied equity volatility. □

Proof of Proposition 3 (Left tail asymmetry)

The adoption payoff $\Pi(\kappa, p, \delta) = p \cdot (\delta \cdot w \cdot \lambda / r) - \kappa$ is bounded above by $\Pi^{\max} = \delta \cdot w \cdot \lambda / r$ (at $p = 1$, $\kappa = 0$) and unbounded below as $\hat{\kappa} \to \infty$. A random variable with a bounded right tail and an unbounded left tail is negatively skewed by construction: the distribution of $\hat{\Phi}$ places strictly more probability mass to the left of its mean than to the right. Failed adoption combined with competitive displacement generates persistent revenue losses with no natural upper bound, making the left tail of the asset value distribution heavier than the right. This concentration of parameter uncertainty in the left tail increases the risk-neutral probability of large negative asset value outcomes, to which deep out-of-the-money put options are most sensitive. High-exposure firms therefore exhibit a larger increase in deep OTM put implied volatility than low-exposure firms following ChatGPT's release. □

Proof of Proposition 4 (Amplification under distress)

We show $\partial^2(CS)/\partial\sigma\,\partial L > 0$ where $L = F \cdot e^{-r\mathfrak{T}}/V_0$. From Part (iii) above, $\partial(CS)/\partial\sigma = n(d_2) \cdot \sqrt{\mathfrak{T}} F e^{-r\mathfrak{T}}/(\mathfrak{T} \cdot D_0)$. Since $d_2 = [\ln(V_0/F) + (r - \sigma^2/2)\mathfrak{T}]/(\sigma\sqrt{\mathfrak{T}})$, we have $\partial d_2/\partial L = -1/(L\sigma\sqrt{\mathfrak{T}}) < 0$. Near the default boundary where $d_2 < 0$, $n'(d_2) = -d_2 \cdot n(d_2) > 0$, so:

$$\frac{\partial^2(CS)}{\partial\sigma\,\partial L} = \frac{\sqrt{\mathfrak{T}}}{\mathfrak{T} \cdot D_0} \cdot n'(d_2) \cdot \frac{\partial d_2}{\partial L} + \text{correction} > 0$$

for sample mean leverage $L \approx 0.40$. Firms closer to the default boundary have smaller d_2, larger $n(d_2)$, and therefore greater spread sensitivity to any volatility perturbation. □

Proof of Proposition 5 (Impact of governance)

From equation (13), $\sigma(\lambda, \Sigma, g) = \sigma_0 + \lambda \cdot \|\Sigma\| \cdot (1 - g)$. By the chain rule:

$$\frac{\partial(CS)}{\partial g} = \left[\frac{n(d_2) \cdot \sqrt{\mathfrak{T}} F \cdot e^{-r\mathfrak{T}}}{\mathfrak{T} \cdot D_0}\right] \cdot [-\lambda \cdot \|\Sigma\|] < 0$$

since all factors are strictly positive and $\lambda > 0$, $\|\Sigma\| > 0$. Higher governance quality g reduces effective parameter uncertainty $\|\Sigma\|(1 - g)$, moderating the risk-neutral volatility increase and therefore the cross-sectional credit spread differential. $\square$

A3 Additional results

Table A3: Core vs. supplemental generative AI exposure decomposition

Notes: This table decomposes total GenAI exposure into core and supplemental components following Eisfeldt, Schubert, and Zhang (2026). Core exposure captures tasks central to an occupation that generative AI can directly substitute; supplemental exposure captures peripheral tasks where generative AI complements rather than replaces workers. The dependent variable is the monthly volume-weighted credit spread in basis points. Column (1) includes core exposure only, column (2) supplemental exposure only, and column (3) both simultaneously. All columns include bond and year-month fixed effects and the full set of controls from Table 2. The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)	(3)
Post × Core Exposure	5.461***		5.679***
	(2.95)		(2.92)
Post × Supplemental Exp		2.719	3.199
		(1.24)	(1.38)
Controls	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Adj. R^2	0.878	0.878	0.878
Observations	308,248	308,248	308,248

Table A4: Volatility channel and left tail risk: Split-sample evidence

Notes: This table reports split-sample evidence supporting Propositions 2 and 3. Columns (1)–(2) split the sample at the median pre-treatment analyst forecast dispersion and re-estimate the baseline specification separately for low-dispersion (below median) and high-dispersion (above median) firms. Columns (3)–(4) split the sample at the median pre-treatment volatility skew (OTM put IV minus ATM IV) and re-estimate the baseline specification separately for low-skew and high-skew firms. The dependent variable is the monthly volume-weighted credit spread in basis points throughout. All regressions include bond and year-month fixed effects and the full set of controls from Table 2. The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)	(3)	(4)
	Low Dispersion	High Dispersion	Low Skew	High Skew
	Panel A: Proposition 2		Panel B: Proposition 3	
Post × GenAI Exp	1.746	14.864***	4.270	9.725***
	(0.89)	(3.18)	(1.64)	(2.97)
Controls	Yes	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Sample	Low Disp.	High Disp.	Low Skew	High Skew
Adj. R^2	0.894	0.864	0.875	0.875
Observations	152,933	152,910	153,224	153,150

Table A5: Credit fragility amplification: Proposition 4

Notes: This table reports split-sample evidence supporting Proposition 4. Each pair of columns splits the sample at the median of the indicated pre-treatment variable and re-estimates the baseline specification separately for below- and above-median firms. The dependent variable is the monthly volume-weighted credit spread in basis points throughout. Columns (1)–(2) split by S&P credit rating (higher values indicate lower ratings); columns (3)–(4) split by equity return volatility; columns (5)–(6) split by distance to default (higher values indicate greater distance from default, i.e., lower credit risk). All regressions include bond and year-month fixed effects and the full set of controls from Table 2. The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)	(3)	(4)	(5)	(6)
	Low	High	Low	High	Low	High
	Rating		*Volatility*		*Dist. to Default*	
Post × GenAI Exp	-2.431	13.302***	-1.339	11.088***	15.106***	-2.878
	(-1.22)	(4.38)	(-0.60)	(3.53)	(4.34)	(-1.52)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.875	0.854	0.887	0.865	0.859	0.881
Observations	156,625	151,199	153,233	154,501	146,497	145,489

Table A6: The role of governance: Proposition 5

Notes: This table reports split-sample evidence supporting Proposition 5. Each pair of columns splits the sample at the median of the indicated pre-treatment variable and re-estimates the baseline specification separately for below- and above-median firms. The dependent variable is the monthly volume-weighted credit spread in basis points throughout. Columns (1)–(2) split by institutional ownership stability (Elyasiani and Jia, 2010). All regressions include bond and year-month fixed effects and the full set of controls from Table 2. The numbers in parentheses are t-statistics. Statistical significance is indicated by *, **, and *** at the 10%, 5%, and 1% levels respectively, using standard errors clustered at the firm level. Monthly data, January 2021 – December 2024.

	(1)	(2)
	Low	High
	Institutional	
Post × GenAI Exp	10.044***	0.826
	(3.15)	(0.34)
Controls	Yes	Yes
Bond FE	Yes	Yes
Year-Month FE	Yes	Yes
Adj. R^2	0.867	0.887
Observations	152,653	149,553